Junos® OS

Junos OS Subscriber Management Feature Guide

Release
13.2

Published: 2013-07-24

Copyright © 2013, Juniper Networks, Inc.

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

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.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Junos® OS Junos OS Subscriber Management Feature Guide
Release 13.2
Copyright © 2013, Juniper Networks, Inc.
All rights reserved.

Revision History
August 2013—Junos OS 13.2

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.

Abbreviated Table of Contents

About This Guide lxvii

Part 1 Managing Access Networks

Chapter 1 Subscriber Access Overview 3

Part 2 AAA Service Framework for Subscriber Access

Chapter 2 Configuring the AAA Service Framework for Subscriber Access 19

Chapter 3 Configuring Address-Assignment Pools for Subscriber Access 167

Chapter 4 Configuring Domain Maps for Subscriber Access 179

Chapter 5 AAA and Remote Subscriber Access Configuration Examples 193

Part 3 DHCP Local Server for Subscriber Access

Chapter 6 DHCP Local Server Overview 199

Chapter 7 Configuring DHCP Local Server 211

Chapter 8 DHCP Local Server Examples 263

Part 4 DHCP Relay Agent for Subscriber Access

Chapter 9 DHCP Relay Agent Overview 275

Chapter 10 Configuring DHCP Relay Agent 285

Chapter 11 DHCP Relay Agent Examples 341

Part 5 PPP for Subscriber Access

Chapter 12 Dynamic Profiles for PPP Overview 363

Chapter 13 Configuring PPP for Subscriber Access 365

Chapter 14 Configuring Subscriber Services for MLPPP Interfaces 373

Part 6 L2TP for Subscriber Access

Chapter 15 L2TP for Subscriber Access Overview 379

Chapter 16 Configuring L2TP for Subscriber Access 401

Part 7 Diameter Base Protocol and Applications for Subscriber Access

Chapter 17 Diameter Base Protocol Overview 455

Chapter 18 Configuring Diameter Base Protocol 471

Copyright © 2013, Juniper Networks, Inc.

Chapter 19 JSRC and Juniper Networks Session Resource Control (SRC) Overview 485
Chapter 20 Configuring JSRC for Subscriber Access 491
Chapter 21 Subscribers on Static Interfaces 495
Chapter 22 Configuring Subscribers over Static Interfaces 499
Chapter 23 PTSP and Juniper Networks Session and Resource Control (SRC) 515
Chapter 24 Configuring the PTSP Application 525
Chapter 25 Configuring Packet-Triggered Subscriber Services 529
Chapter 26 Gx-Plus Overview 541
Chapter 27 Configuring Gx-Plus 549

Part 8 Mobile IP Access
Chapter 28 Mobile IP Overview 557
Chapter 29 Configuring Mobile IP 569

Part 9 ATM for Subscriber Access
Chapter 30 ATM for Subscriber Access Overview 581
Chapter 31 Configuring ATM for Subscriber Access 587
Chapter 32 ATM for Subscriber Access Examples 593

Part 10 Dynamic Profiles for Access and Services
Chapter 33 Dynamic Profiles Overview 637
Chapter 34 Configuring Dynamic Profiles 673
Chapter 35 Configuring Dynamic Access and Access-Internal Routes for Subscriber Management 687
Chapter 36 Dynamic Profile Examples 695

Part 11 Dynamic VLANs
Chapter 37 Dynamic VLAN Overview 701
Chapter 38 Configuring Dynamic VLANs 713
Chapter 39 Dynamic VLAN Examples 751

Part 12 Subscriber Interfaces
Chapter 40 Subscriber Interface Overview 757
Chapter 41 Configuring Subscriber Interfaces for Dynamic Profiles 765
Chapter 42 Subscriber Interface Examples 779
Chapter 43 Subscriber Interfaces over Aggregated Ethernet Overview 815
Chapter 44 Configuring Subscriber Interfaces over Aggregated Ethernet 823
Chapter 45 Subscriber Interfaces over Aggregated Ethernet Examples 835
Chapter 46 Dynamic PPPoE Subscriber Interfaces Overview 883
Chapter 47 Configuring Dynamic PPPoE Subscriber Interfaces 899

 Copyright © 2013, Juniper Networks, Inc.

Chapter 48 Dynamic PPPoE Subscriber Interfaces Examples 917
Chapter 49 MPLS Pseudowire Subscriber Interfaces Overview 923
Chapter 50 Configuring MPLS Pseudowire Subscriber Interfaces 933

Part 13 Class of Service for Subscriber Access

Chapter 51 Dynamic CoS for Subscriber Access Overview 949
Chapter 52 Configuration Summary of Dynamic CoS for Subscriber Access 963
Chapter 53 Configuring Dynamic Shaping and Scheduling for Subscriber Access . . . 971
Chapter 54 RADIUS and Dynamic CoS Overview 989
Chapter 55 Configuring RADIUS for Dynamic CoS 999
Chapter 56 Interface Solutions for Dynamic CoS Overview 1013
Chapter 57 Configuring Interface Solutions for Dynamic CoS 1023
Chapter 58 Dynamic CoS for Subscriber Access Examples 1031
Chapter 59 Bandwidth Management for Dynamic CoS Overview 1069
Chapter 60 Configuring Bandwidth Management Parameters for Dynamic CoS . . . 1089
Chapter 61 Bandwidth Management for Dynamic CoS Examples 1111

Part 14 Subscriber Service Activation

Chapter 62 Subscriber Service Activation Overview 1121

Part 15 Dynamic Firewall Filters, Service Sets and HTTP Redirect for Subscriber Access

Chapter 63 Dynamic Firewall Filters and Service Sets Overview 1131
Chapter 64 Parameterized Filters 1153
Chapter 65 Configuring Filters for Dynamic Profiles 1169
Chapter 66 Configuring Fast Update Filters 1183
Chapter 67 Configuring Service Sets in Dynamic Profiles 1197
Chapter 68 Firewall Filter Examples 1199
Chapter 69 Redirecting HTTP Requests Overview 1225
Chapter 70 Configuring HTTP Redirect 1229
Chapter 71 HTTP Redirect Examples 1233

Part 16 Subscriber Secure Policy Traffic Mirroring

Chapter 72 Subscriber Secure Policy Traffic Mirroring Overview 1245
Chapter 73 Overview of Subscriber Secure Policy Using RADIUS 1247
Chapter 74 Configuring RADIUS-Initiated Subscriber Secure Policy Traffic Mirroring 1261
Chapter 75 Overview of Subscriber Secure Policy Using DTCP 1267
Chapter 76 Configuring DTCP-Initiated Subscriber Secure Policy Mirroring 1281

Copyright © 2013, Juniper Networks, Inc.

Chapter 77 DTCP Messages Used for DTCP-Initiated Subscriber Secure Policy Mirroring 1291
Chapter 78 Reporting Intercept-Related Events for Subscriber Secure Policy 1305

Part 17 Protocols for Subscriber Access

Chapter 79 ANCP Overview 1311
Chapter 80 Configuring ANCP 1331
Chapter 81 Dynamic IGMP Configuration Overview 1371
Chapter 82 Dynamic MLD Configuration Overview 1373
Chapter 83 Dynamic Router Advertisement Overview 1375

Part 18 Subscriber Access Examples

Chapter 84 Service Profile Examples 1379

Part 19 Complete Configuration Statement Hierarchy and Summary of Statements for Subscriber Access

Chapter 85 Subscriber Access Statement Hierarchy 1385
Chapter 86 Subscriber Access Configuration Statements 1413

Part 20 Index

Index 2123
Index of Statements and Commands 2167

 Copyright © 2013, Juniper Networks, Inc.

Table of Contents

About This Guide lxvii

JUNOS Documentation and Release Notes lxvii
Objectives lxviii
Audience lxviii
Supported Routing Platforms lxix
Using the Indexes lxix
Using the Examples in This Manual lxix
Merging a Full Example lxix
Merging a Snippet lxx
Documentation Conventions lxx
Documentation Feedback lxxii
Requesting Technical Support lxxii
Self-Help Online Tools and Resources lxxiii
Opening a Case with JTAC lxxiii

Part 1 Managing Access Networks

Chapter 1 **Subscriber Access Overview 3**

Subscriber Access Overview 3
Subscriber Access Terms and Acronyms 4
Subscriber Access Environment 4
Relationship Between Subscribers and Interfaces in an Access Network 5
Subscriber Access Support Considerations 5
Platform Support 5
Interface Support 6
DPC Support 6
Routing Engine Support 6
Logical System Support 6
Subscriber Access Licensing Overview 6
Subscriber Management Unified ISSU Support 7
Unified ISSU Support for DHCP Access Model 7
Unified ISSU Support for PPPoE Access Model 7
Verifying and Monitoring Subscriber Management Unified ISSU State 8
Subscriber Access Operation Flow Using DHCP Relay 8

Subscriber Activation and Service Management in an Access Network 9
Components of a Dynamic Profile 10
Router Predefined Variables Used by Dynamic Profiles 10
Configuring Subscriber Access 10
Collecting Subscriber Access Logs Before Contacting Juniper Technical Support 13
Compressing Troubleshooting Logs from /var/logs to Send to Juniper Technical Support 15

Part 2 AAA Service Framework for Subscriber Access

Chapter 2 **Configuring the AAA Service Framework for Subscriber Access 19**

AAA Service Framework Overview 21
Configuring Router or Switch Interaction with RADIUS Servers 23
Configuring Authentication and Accounting Parameters for Subscriber Access 24
Specifying the Authentication and Accounting Methods for Subscriber Access 24
RADIUS Acct-On and Acct-Off Messages 25
RADIUS Accounting Statistics for Subscriber Access Overview 26
Understanding RADIUS Accounting Duplicate Reporting 28
Layer 3 Wholesale Scenarios 28
Other Scenarios 29
Configuring Per-Subscriber Session Accounting 29
Retaining Authentication and Accounting Information During Session Startup 31
Preservation of RADIUS Accounting Information During an Accounting Server Outage 32
Configuring Per-Service Session Accounting 35
Configuring RADIUS Server Parameters for Subscriber Access 36
Specifying RADIUS Authentication and Accounting Servers for Subscriber Access 37
Configuring Back-up Options for RADIUS Accounting 38
Forcing the Router to Contact the Accounting Server Immediately 39
Monitoring Pending RADIUS Accounting Stop Messages 39
RADIUS Server Options for Subscriber Access 41
Global RADIUS Options for Subscriber Access 44
DNS Address Assignment Precedence 44
Manual Configuration of the NAS-Port-ID RADIUS Attribute 44
Manual Configuration of the NAS-Port-Type RADIUS Attribute 45
Configuring RADIUS Server Options for Subscriber Access 47
Configuring RADIUS Options for Subscriber Access Globally 50
Subscriber Access Interface Description Storage and Reporting Through RADIUS Overview 50
Interface Description Precedence 51
Example: Reporting Interface Descriptions on Non-Underlying Logical Interfaces 51
Reporting Interface Descriptions on Underlying Logical Interfaces 52
Interface Descriptions on Aggregated Ethernet Physical Interfaces 53

 Copyright © 2013, Juniper Networks, Inc.

Interface Descriptions on a Combination of Dynamic and Static Interfaces . . . 53
Example: Reporting Interface Descriptions on Dynamic VLANs . . . 53
Configuring a NAS-Port-ID with Additional Options . . . 54
Configuring a Calling-Station-ID with Additional Attributes . . . 56
Configuring How RADIUS Attributes Are Used for Subscriber Access . . . 58
RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview . . . 63
NAS-Port-Type RADIUS Attribute . . . 63
NAS-Port RADIUS Attribute . . . 64
NAS-Port Options Configuration and Subscriber Network Access Models . . . 64
NAS-Port Options Definition . . . 64
Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN . . . 65
Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN . . . 66
Configuring the RADIUS NAS-Port-Type per Physical Interface . . . 68
Configuring the RADIUS NAS-Port-Type per VLAN . . . 69
Configuring the RADIUS NAS-Port-Type per Stacked VLAN . . . 70
Configuring the RADIUS NAS-Port Extended Format per Physical Interface . . . 72
Configuring the RADIUS NAS-Port Extended Format per VLAN . . . 74
Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN . . . 75
Configuring the RADIUS NAS-Port Extended Format for ATM Interfaces . . . 78
Centrally Configured Opaque DHCP Options . . . 80
Data Flow for RADIUS-Sourced DHCP Options . . . 82
Multiple VSA 26-55 Instances Configuration . . . 83
DHCP Options That Cannot Be Centrally Configured . . . 83
Monitoring DHCP Options Configured on RADIUS Servers . . . 84
Using RADIUS Dynamic Requests for Subscriber Access Management . . . 86
Dynamic Service Activation During Login Overview . . . 86
RADIUS-Initiated Change of Authorization (CoA) Overview . . . 87
CoA Messages . . . 87
Qualifications for Change of Authorization . . . 87
Message Exchange . . . 88
RADIUS-Initiated Disconnect Overview . . . 88
Disconnect Messages . . . 89
Qualifications for Disconnect . . . 89
Message Exchange . . . 89
Configuring RADIUS-Initiated Dynamic Request Support . . . 90
Verifying and Managing the RADIUS Dynamic-Request Feature . . . 90
RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework . . . 90
RADIUS IETF Attributes Supported by the AAA Service Framework . . . 91
Juniper Networks VSAs Supported by the AAA Service Framework . . . 98
DSL Forum Vendor-Specific Attributes . . . 108
AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS . . . 110

AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS . . . 115
DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS . . . 119
Error-Cause Codes (RADIUS Attribute 101) for Dynamic Requests . . . 120
Mapping Application Terminate Reasons and RADIUS Terminate Codes . . . 121
AAA Terminate Reasons . . . 123
DHCP Terminate Reasons . . . 124
L2TP Terminate Reasons . . . 124
PPP Terminate Reasons . . . 140
Configuring Custom Terminate Reason Mappings . . . 148
Configuring an Access Profile for Subscriber Management . . . 149
Attaching Access Profiles . . . 150
DNS Name Server Address Overview . . . 150
Configuring DNS Name Server Addresses for Subscriber Management . . . 151
Understanding Session Options for Subscriber Access . . . 153
Configuring Subscriber Session Options . . . 154
Removing Inactive Dynamic Subscriber VLANs . . . 155
AAA Configuration Testing and Troubleshooting . . . 156
Testing a Subscriber AAA Configuration . . . 156
Tracing Subscriber Management Database Operations for Subscriber Access . . . 158
Configuring the Subscriber Management Database Trace Log Filename . . . 159
Configuring the Number and Size of Subscriber Management Database Log Files . . . 159
Configuring Access to the Subscriber Management Database Log File . . . 160
Configuring a Regular Expression for Subscriber Management Database Messages to Be Logged . . . 160
Configuring the Subscriber Management Database Tracing Flags . . . 161
Tracing Subscriber Management Session Database Replication Operations for Subscriber Access . . . 161
Configuring the Number and Size of Subscriber Management Session Database Replication Log Files . . . 162
Configuring a Regular Expression for Subscriber Management Session Database Replication Messages to Be Logged . . . 163
Configuring Access to the Subscriber Management Session Database Replication Log File . . . 163
Configuring the Subscriber Management Session Database Replication Trace Log Filename . . . 164
Configuring the Subscriber Management Session Database Replication Tracing Flags . . . 164
Verifying and Managing Subscriber AAA Information . . . 164

Chapter 3 **Configuring Address-Assignment Pools for Subscriber Access . . . 167**

Address-Assignment Pools Overview . . . 167
Configuring Address-Assignment Pools . . . 168
Configuring an Address-Assignment Pool Name and Addresses . . . 169
Configuring a Named Address Range for Dynamic Address Assignment . . . 169
Configuring Address-Assignment Pool Linking . . . 170

 Copyright © 2013, Juniper Networks, Inc.

Configuring Static Address Assignment 171
Configuring an Address-Assignment Pool for Router Advertisement 171
Configuring Address-Assignment Pool Usage Threshold Traps 172
Configuring DHCP Client-Specific Attributes 172
DHCP Attributes for Address-Assignment Pools 173
Address-Assignment Pools Licensing Requirements 175
Tracing General Authentication Service Processes 175
Configuring the General Authentication Service Processes Trace Log Filename 176
Configuring the Number and Size of General Authentication Service Processes Log Files 176
Configuring Access to the Log File 177
Configuring a Regular Expression for Lines to Be Logged 177
Configuring the Trace Operation 178

Chapter 4 Configuring Domain Maps for Subscriber Access 179

Domain Mapping Overview 180
Default Domain Map 181
Configuring a Domain Map 181
Specifying an Access Profile in a Domain Map 182
Specifying a Dynamic Profile in a Domain Map 183
Specifying an Address Pool in a Domain Map 184
Specifying an AAA Logical System/Routing Instance in a Domain Map 184
Specifying a Target Logical System/Routing Instance in a Domain Map 186
Configuring Domain Name Usage for Domain Maps 187
Specifying Domain Name Delimiters 187
Specifying the Parsing Direction for Domain Names 188
Enabling Domain Name Stripping 189
Specifying a Tunnel Profile in a Domain Map 189
Specifying a Tunnel Switch Profile in a Domain Map 190
Configuring PADN Parameters for a Domain Map 190
Verifying and Managing Domain Map Configuration 191

Chapter 5 AAA and Remote Subscriber Access Configuration Examples 193

Example: Configuring RADIUS-Based Subscriber Authentication and Accounting 193
Example: Configuring an Address-Assignment Pool 195

Part 3 DHCP Local Server for Subscriber Access

Chapter 6 DHCP Local Server Overview 199

Extended DHCP Local Server Overview 200
Interaction Among the DHCP Client, Extended DHCP Local Server, and Address-Assignment Pools 200
Providing DHCP Client Configuration Information 201
Minimal Configuration for Clients 202
DHCP Local Server and Address-Assignment Pools 202
DHCP Liveness Detection 203
DHCPv6 Local Server Overview 204
DHCP Local Server Handling of Client Information Request Messages 205

Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview 207
Multiple DHCP Subscribers Sharing the Same VLAN Logical Interface 207
Primary Dynamic Profile 208
Port Number Requirements for DHCP Firewall Filters 208

Chapter 7 **Configuring DHCP Local Server 211**

DHCP Duplicate Client Differentiation Using Client Subinterface Overview 212
Guidelines for Configuring Support for DHCP Duplicate Clients 213
Configuring DHCP Duplicate Client Support 214
Using External AAA Authentication Services with DHCP 214
Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use 215
Grouping Interfaces with Common DHCP Configurations 217
Guidelines for Configuring Interface Ranges 218
Group-Specific DHCP Local Server Options 219
Overriding Default DHCP Local Server Configuration Settings 220
Specifying the Maximum Number of DHCP Clients Per Interface 221
DHCP Access, Access-Internal, and Destination Route Suppression 222
Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default 223
Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server 223
DHCP Auto Logout Overview 225
Auto Logout Overview 225
How DHCP Identifies and Releases Clients 225
Option 60 and Option 82 Requirements 226
Automatically Logging Out DHCP Clients 226
Enabling Processing of Client Information Requests 227
Enabling DHCPv6 Rapid Commit Support 228
Deleting DHCP Local Server and DHCP Relay Override Settings 229
Use of DHCP Option 50 and DHCPv6 IA_NA Option to Request a Specific IP Address 230
Multiple Address Assignment for DHCPv6 Clients 230
Multiple Address Assignment Using Local Address Pools or RADIUS 230
Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment . . 231
Specifying the Delegated Address Pool for IPv6 Prefix Assignment 231
Subscriber Binding Retention During Interface Delete Events 232
Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events 232
Verifying and Managing the DHCP Maintain Subscribers Feature 233
Clearing DHCP Bindings for Subscriber Access 234
Configuring the Router to Strictly Enforce the Subscriber Scaling License 235
Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces 236
Attaching a Dynamic Profile to All DHCP Subscriber or All DHCP Client Interfaces 236
Attaching a Dynamic Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces 237
Configuring Passwords for Usernames 238

 Copyright © 2013, Juniper Networks, Inc.

Creating Unique Usernames for DHCP Clients 239
Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients 241
Default Client/Server Interaction 241
Dynamic Client/Server Interaction for DHCPv4 242
Dynamic Client/Server Interaction for DHCPv6 242
Dynamic Configuration Options 243
Configuring Extended DHCP Local Server Dynamic Client Reconfiguration 245
Configuring Dynamic Reconfiguration Attempts for DHCP Clients 246
Configuring Deletion of the Client When Dynamic Reconfiguration Fails 247
Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect 247
Configuring a Token for DHCP Local Server Authentication 248
Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings . . 248
Preventing Binding of Clients That Do Not Support Reconfigure Messages 249
DHCP Liveness Detection Overview 250
Configuring Detection of DHCP Local Server Client Connectivity 251
Verifying and Managing DHCP Local Server Configuration 253
Verifying and Managing DHCPv6 Local Server Configuration 253
Graceful Routing Engine Switchover 254
Tracing Extended DHCP Operations 255
Configuring the Extended DHCP Log Filename 256
Configuring the Number and Size of Extended DHCP Log Files 257
Configuring Access to the Extended DHCP Log File 257
Configuring a Regular Expression for Extended DHCP Messages to Be Logged 258
Configuring the Extended DHCP Tracing Flags 258
Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged 259
Tracing Extended DHCP Operations for Specific Interfaces 260

Chapter 8 DHCP Local Server Examples 263

Example: Minimum Extended DHCP Local Server Configuration 263
Example: Extended DHCP Local Server Configuration with Optional Pool Matching 263
Example: Extended DHCPv6 Local Server Configuration 264
Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine 265
Example: Configuring Group Liveness Detection for DHCP Local Server Clients 269

Part 4 DHCP Relay Agent for Subscriber Access

Chapter 9 DHCP Relay Agent Overview 275

Extended DHCP Relay Agent Overview 276
Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers 278
DHCP Liveness Detection 279
DHCPv6 Relay Agent Overview 279
DHCP Relay Proxy Overview 280
Interaction Among DHCP Relay Proxy, DHCP Client, and DHCP Servers . . . 281

Copyright © 2013, Juniper Networks, Inc.

Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview 282
Multiple DHCP Subscribers Sharing the Same VLAN Logical Interface 282
Primary Dynamic Profile 283

Chapter 10 **Configuring DHCP Relay Agent 285**

DHCP Duplicate Client Differentiation Using Client Subinterface Overview 286
Guidelines for Configuring Support for DHCP Duplicate Clients 287
Configuring DHCP Duplicate Client Support 288
Using External AAA Authentication Services with DHCP 288
Grouping Interfaces with Common DHCP Configurations 290
Guidelines for Configuring Interface Ranges 291
Group-Specific DHCP Relay Options 292
Overriding the Default DHCP Relay Configuration Settings 293
Overwriting giaddr Information 294
Replacing the DHCP Relay Request and Release Packet Source Address 295
Overriding Option 82 Information 295
Using Layer 2 Unicast Transmission for DHCP Packets 296
Trusting Option 82 Information 296
DHCP Access, Access-Internal, and Destination Route Suppression 297
Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default 298
Specifying the Maximum Number of DHCP Clients Per Interface 298
DHCP Snooping Support 299
Configuring DHCP Snooping for DHCP Relay Agent 301
Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent 301
Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent 306
DHCP Auto Logout Overview 308
Auto Logout Overview 308
How DHCP Identifies and Releases Clients 309
Option 60 and Option 82 Requirements 309
DHCP Relay Agent Option 82 Value for Auto Logout 310
Automatically Logging Out DHCP Clients 311
Sending Release Messages When Clients Are Deleted 312
Subscriber Binding Retention During Interface Delete Events 313
Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events 313
Verifying and Managing the DHCP Maintain Subscribers Feature 314
Clearing DHCP Bindings for Subscriber Access 315
DHCP Liveness Detection Overview 316
Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity 318
Disabling DHCP Relay 319
Disabling Automatic Binding of Stray DHCP Requests 320
DHCP Options and Selective Traffic Processing Overview 321
Using DHCP Option Information to Selectively Process DHCP Client Traffic . . . 323

 Copyright © 2013, Juniper Networks, Inc.

Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings 325
Enabling and Disabling Insertion of Option 82 Information 325
Configuring Agent Circuit ID Information 325
Configuring an Option 82 Prefix 326
Using a Textual Description in Option 82 327
Configuring Server Groups 328
Configuring Active Server Groups 328
Enabling DHCP Relay Proxy Mode 329
Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets 329
Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces 330
Attaching a Dynamic Profile to All DHCP Subscriber or All DHCP Client Interfaces 331
Attaching a Dynamic Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces 331
Verifying and Managing DHCP Relay Configuration 332
Verifying and Managing DHCPv6 Relay Configuration 333
Tracing Extended DHCP Operations 333
Configuring the Extended DHCP Log Filename 335
Configuring the Number and Size of Extended DHCP Log Files 335
Configuring Access to the Extended DHCP Log File 336
Configuring a Regular Expression for Extended DHCP Messages to Be Logged 336
Configuring the Extended DHCP Tracing Flags 337
Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged 337
Tracing Extended DHCP Operations for Specific Interfaces 338
Disabling DHCP Relay 339

Chapter 11 DHCP Relay Agent Examples 341

Example: Minimum DHCP Relay Agent Configuration 341
Example: DHCP Relay Agent Configuration with Multiple Clients and Servers . . 342
Example: Configuring DHCP Snooping Support for DHCP Relay Agent 343
Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent 345
Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings 349
Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing 353
Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients 357

Part 5 PPP for Subscriber Access

Chapter 12 Dynamic Profiles for PPP Overview 363

Dynamic Profiles for PPP Subscriber Interfaces Overview 363

Chapter 13 Configuring PPP for Subscriber Access 365

Configuring Dynamic Authentication for PPP Subscribers 365
Controlling the Negotiation Order of PPP Authentication Protocols 367

Modifying the CHAP Challenge Length 369
Attaching Dynamic Profiles to Static PPP Subscriber Interfaces 370
Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests 371
How PPP Fast Keepalive Processing Works 371
Statistics Display for PPP Fast Keepalive 372
Effect of Changing the Forwarding Class Configuration 372
Verifying and Managing PPP Configuration for Subscriber Management 372

Chapter 14 **Configuring Subscriber Services for MLPPP Interfaces 373**

Dynamic PPP Subscriber Services for Static MLPPP Interfaces 373
Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces 374
Configuring PPP Subscriber Services for MLPPP Bundles 374
Enabling PPP Subscriber Services for Static Non-Ethernet Interfaces 375
Attaching Dynamic Profiles to MLPPP Bundles 375

Part 6 L2TP for Subscriber Access

Chapter 15 **L2TP for Subscriber Access Overview 379**

L2TP for Subscriber Access Overview 379
L2TP Terminology 382
L2TP Implementation 383
Sequence of Events on the LAC 383
Sequence of Events on the LNS 384
LAC Tunnel Selection Overview 384
Tunnel Selection Failover Between Preference Levels 385
Tunnel Selection Failover Within a Preference Level 386
Tunnel Selection and Maximum Sessions per Tunnel 387
Tunnel Selection with Weighted Load Balancing 388
L2TP Failover and Peer Resynchronization 388
L2TP and Graceful Routing Engine Switchover 389
IP Packet Fragment Reassembly for L2TP Overview 390
Transmission of Tx Connect-Speed and Rx Connect-Speed AVPs from LAC to LNS 391
L2TP Tunnel Switching Overview 392
Application of Tunnel Switch Profiles 394
Termination of Tunnel-Switched Sessions on the LTS 395
Tunnel Switching Actions for L2TP AVPs at the Switching Boundary 396
LAC Interoperation with Third-Party LNS Devices 398

Chapter 16 **Configuring L2TP for Subscriber Access 401**

Configuring an L2TP LAC 402
Configuring a Tunnel Profile for Subscriber Access 403
Configuring the L2TP LAC Tunnel Selection Parameters 406
Configuring LAC Tunnel Selection Failover Within a Preference Level 406
Configuring Weighted Load Balancing for LAC Tunnel Sessions 407
Preventing the LAC from Sending Calling Number AVP 22 to the LNS 407
Configuring the Method to Set the LAC Connection Speeds to the LNS 408

 Copyright © 2013, Juniper Networks, Inc.

Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal . . . 409
Preventing the LAC From Negotiating L2TP Failover Protocol . . . 409
Setting the Format for the Tunnel Name . . . 410
Configuring the Number of L2TP Control Message Retransmissions . . . 411
Setting the L2TP Tunnel Idle Timeout . . . 412
Setting the L2TP Receive Window Size . . . 412
Setting the L2TP Destruct Timeout . . . 413
Configuring the LAC to Ignore Address and Port Changes Requested by the LNS . . . 413
Configuring the L2TP Destination Lockout Timeout . . . 414
Removing an L2TP Destination from the Destination Lockout List . . . 415
Enabling Tunnel and Global Counters for SNMP Statistics Collection . . . 416
Configuring an L2TP LNS with Inline Service Interfaces . . . 417
Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile . . . 418
Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface . . . 419
Configuring IP Inline Reassembly for L2TP . . . 421
Configuring an L2TP Access Profile on the LNS . . . 422
Configuring a AAA Local Access Profile on the LNS . . . 424
Configuring an Address-Assignment Pool for L2TP LNS with Inline Services . . . 424
Configuring the L2TP LNS Peer Interface . . . 426
Enabling Inline Service Interfaces . . . 427
Configuring an Inline Service Interface for L2TP LNS . . . 427
Configuring Options for the LNS Inline Services Logical Interface . . . 428
Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces . . . 429
Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions . . . 430
Configuring a Dynamic Profile for Dynamic LNS Sessions . . . 431
Configuring L2TP Tunnel Switching . . . 432
Tracing L2TP Operations for Subscriber Access . . . 434
Configuring the L2TP Trace Log Filename . . . 435
Configuring the Number and Size of L2TP Log Files . . . 435
Configuring Access to the L2TP Log File . . . 436
Configuring a Regular Expression for L2TP Messages to Be Logged . . . 436
Configuring the L2TP Tracing Flags . . . 437
Configuring the Severity Level to Filter Which L2TP Messages Are Logged . . . 437
Verifying and Managing L2TP for Subscriber Access . . . 438
Testing L2TP Tunnel Configurations from the LAC . . . 439
Example: Configuring an L2TP LNS . . . 440

Part 7 Diameter Base Protocol and Applications for Subscriber Access

Chapter 17 **Diameter Base Protocol Overview . . . 455**

Diameter Base Protocol Overview . . . 455
Messages Used by Diameter Applications . . . 457
Diameter AVPs and Diameter Applications . . . 462

Chapter 18 **Configuring Diameter Base Protocol . . . 471**

Configuring Diameter . . . 471
Configuring the Origin Attributes of the Diameter Instance . . . 472

Configuring Diameter Peers 473
Configuring Diameter Network Elements 473
Configuring the Diameter Transport 474
Tracing Diameter Base Protocol Processes for Subscriber Access 475
Configuring the Diameter Base Protocol Trace Log Filename 476
Configuring the Number and Size of Diameter Base Protocol Log Files 476
Configuring Access to the Diameter Base Protocol Log File 477
Configuring a Regular Expression for Diameter Base Protocol Messages to Be Logged 477
Configuring the Diameter Base Protocol Tracing Flags 478
Configuring the Severity Level to Filter Which Diameter Base Protocol Messages Are Logged 478
Verifying Diameter Node, Instance, and Route Information 478
Verifying and Managing Diameter Function Information 479
Verifying and Managing Diameter Peer Information 480
Verifying Diameter Network Element Information 481
Troubleshooting Diameter Network Configuration 482
Troubleshooting Diameter Network Connectivity 482

Chapter 19 JSRC and Juniper Networks Session Resource Control (SRC) Overview 485

Juniper Networks Session and Resource Control (SRC) and JSRC Overview . . . 485
Hardware Requirements for JSRC for Subscriber Access 486
Service Accounting with JSRC 486
Understanding JSRC-SAE Interactions 488
Subscriber Login 488
Subscriber Service Activation and Deactivation 489
Subscriber Resynchronization 489
Subscriber Session Terminated by the SAE 489
Statistics Collection and Reporting per Service Rule 490
Subscriber Logout 490

Chapter 20 Configuring JSRC for Subscriber Access 491

Configuring JSRC 491
Configuring the JSRC Partition 492
Assigning a Partition to JSRC 493
Authorizing Subscribers with JSRC 493
Provisioning Subscribers with JSRC 494
Configuring Service Accounting with JSRC 494

Chapter 21 Subscribers on Static Interfaces 495

Subscribers on Static Interfaces Overview 495

Chapter 22 Configuring Subscribers over Static Interfaces 499

Configuring Subscribers over Static Interfaces 500
Specifying the Static Subscriber Global Access Profile 501
Specifying the Static Subscriber Global Dynamic Profile 502
Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers 502
Configuring the Static Subscriber Global Authentication Password 503
Configuring the Static Subscriber Global Username 503

Copyright © 2013, Juniper Networks, Inc.

Creating a Static Subscriber Group 504
Specifying the Static Subscriber Group Access Profile 505
Specifying the Static Subscriber Group Dynamic Profile 505
Enabling Multiple Subscribers on a VLAN Logical Interface for a Static Subscriber Group 506
Configuring the Static Subscriber Group Authentication Password 506
Configuring the Static Subscriber Group Username 507
Forcing a Static Subscriber to Be Logged Out 508
Resetting the State of an Interface for Static Subscriber Login 508
Forcing a Group of Static Subscribers to Be Logged Out 508
Resetting the State of an Interface Group for Static Subscriber Login 508
Verifying Information about Subscriber Sessions on Static Interfaces 509
Tracing Static Subscriber Operations 509
Configuring the Static Subscribers Trace Log Filename 510
Configuring the Number and Size of Static Subscribers Log Files 510
Configuring Access to the Static Subscribers Log File 511
Configuring a Regular Expression for Static Subscriber Messages to Be Logged 511
Configuring the Static Subscribers Tracing Flags 512
Configuring the Severity Level to Filter Which Static Subscriber Messages Are Logged 512
Example: Configuring Static Subscribers for Subscriber Access 512

Chapter 23 PTSP and Juniper Networks Session and Resource Control (SRC) 515

PTSP Overview 515
Hardware Requirements for PTSP for Subscriber Access 516
Juniper Networks Session and Resource Control (SRC) and PTSP Overview . . . 516
Understanding PTSP-SAE Interactions 517
Packet-Triggered Subscribers Services Overview 518
Subscriber Identification Method for PTSP Partition 520
PTSP Services on Aggregated and Redundant Services PICs 520
Understanding the Subscriber Profiles for Client Sessions per PSTP Partition . . 522

Chapter 24 Configuring the PTSP Application 525

Configuring the PTSP Application 525
Configuring the PTSP Partition 526
Assigning the PTSP Partition 526
Tracing Packet-Triggered Subscriber Operations 527
Configuring the Packet-Triggered Subscribers Trace Log Filename 528
Configuring the Size of Packet-Triggered Subscribers Log Files 528
Configuring the Packet-Triggered Subscribers Tracing Flags 528

Chapter 25 Configuring Packet-Triggered Subscriber Services 529

Configuring PTSP 529
Configuring the Multiservices DPC for PTSP 530
Enabling the PTSP Service Package on the Multiservices DPC 530
Configuring Services Interface for PTSP 530
Configuring PTSP Service Rules 531
Configuring Static PTSP Rules 531
Configuring PTSP Rule Sets 533

Copyright © 2013, Juniper Networks, Inc.

Configuring PTSP Service Sets 534
Configuring the PTSP Forwarding Instance 534
Configuring a Statistics Profile for PTSP 536
Configuring the File Properties for Statistics Data Output 537
Configuring the Profile Properties for Statistics Data Output 537
Configuring the Record Type for Statistics Data 538
Tracing PTSP Operations 538
Verifying and Managing PTSP Configuration 539

Chapter 26 **Gx-Plus Overview 541**

Gx-Plus for Provisioning Subscribers Overview 541
Understanding Gx-Plus Interactions Between the Router and the PCRF 543
Subscriber Login 543
Fault Tolerance and Event Notification 545
PCRF-Generated Discovery 546
Subscriber Accounting 547
Subscriber Audit 547
Subscriber Logout 547

Chapter 27 **Configuring Gx-Plus 549**

Configuring Gx-Plus 549
Configuring the Gx-Plus Partition 550
Configuring Gx-Plus Global Attributes 551
Provisioning Subscribers with Gx-Plus 552
Disabling PCRF Control of a Subscriber Session 552

Part 8 Mobile IP Access

Chapter 28 **Mobile IP Overview 557**

Mobile IP Home Agent Elements and Behavior 557
Mobile IP Registration 560
Home Address Assignment 560
Authentication 560
Reauthentication 561
AAA Authentication 561
Local Authentication 562
Accounting 563
Mobile IP Routing and Forwarding 564
Mobile IP in the WiMAX Environment 565

Chapter 29 **Configuring Mobile IP 569**

Configuring Mobile IP 569
Configuring the Mobile IP Authentication Method 570
Configuring the Mobile IP Home Agent 570
Configuring the Local Authentication Attributes for the Mobile Node 571
Configuring Accounting for Mobile IP Subscribers 572
Configuring Dynamic Home Assignment for the Mobile Node 572
Configuring the Access Type for Mobile IP 573
Tracing Mobile IP Operations for Subscriber Access 573
Configuring the Mobile IP Trace Log Filename 575

 Copyright © 2013, Juniper Networks, Inc.

Configuring the Number and Size of Mobile IP Log Files 575
Configuring Access to the Mobile IP Log File 576
Configuring a Regular Expression for Mobile IP Messages to Be Logged 576
Configuring the Mobile IP Tracing Flags 576
Configuring the Severity Level to Filter Which Mobile IP Messages Are Logged 577

Part 9 ATM for Subscriber Access

Chapter 30 **ATM for Subscriber Access Overview 581**

ATM for Subscriber Access Overview 581
Supported Configurations for ATM Subscriber Access 581
PPP-over-Ethernet-over-ATM Configurations 582
Routed IP-over-ATM Configurations 582
Bridged IP-over-Ethernet-over-ATM Configurations 583
PPP-over-ATM Configurations 583
Configuration and Encapsulation Types for ATM Subscriber Access 584
ATM for Subscriber Access Encapsulation Types Overview 585

Chapter 31 **Configuring ATM for Subscriber Access 587**

Configuring ATM for Subscriber Access 587
Guidelines for Configuring ATM for Subscriber Access 589
Verifying and Managing ATM Configurations for Subscriber Access 590

Chapter 32 **ATM for Subscriber Access Examples 593**

Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM 593
Example: Configuring a Static PPPoE Subscriber Interface over ATM 602
Example: Configuring a Static Subscriber Interface for IP Access over ATM 610
Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM 616
Example: Configuring a Static PPP Subscriber Interface over ATM 623

Part 10 Dynamic Profiles for Access and Services

Chapter 33 **Dynamic Profiles Overview 637**

Dynamic Profiles Overview 638
Dynamic Profile Interface Support 638
What Dynamic Profiles Do 638
How Dynamic Profiles Work 639
Dynamic Profile Version Creation 639
Dynamic Profile Semantic Checking 640
Dynamic Variables Overview 640
How Dynamic Variables Work 641
Default Values for Predefined Variables 641
Unique Identifier (UID) for Parameterized Filters 641
Junos OS Predefined Variables 642
Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs 660
User-Defined Variables 666
Variable Expressions Overview 667

Copyright © 2013, Juniper Networks, Inc.

Access Profiles and Service Profiles Overview 670
Functionality Supported in Access and Service Profiles 670

Chapter 34 **Configuring Dynamic Profiles 673**

Configuring a Basic Dynamic Profile 673
Configuring Predefined Dynamic Variables in Dynamic Profiles 675
Configuring Default Values for Predefined Variables in a Dynamic Profile 675
Configuring User-Defined Dynamic Variables in Dynamic Profiles 676
Configuring Variable Expressions in Dynamic Profiles 678
Configuring a Dynamic Profile for Client Access 679
Configuring a Dynamic Profile for Various Levels of Services 681
Configuring Unique Identifiers for Parameterized Filters in Dynamic Profiles . . . 682
Enabling Dynamic Profiles to use Multiple Versions 683
Modifying Dynamic Profiles with Versioning Disabled 684

Chapter 35 **Configuring Dynamic Access and Access-Internal Routes for Subscriber Management 687**

Access and Access-Internal Routes for Subscriber Management 687
Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover 688
Benefits of Delaying Removal of Access Routes and Access-Internal Routes 689
Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes 689
Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes 689
Configuring Dynamic Access Routes for Subscriber Management 690
Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management 692
Configuring Dynamic Access-Internal Routes for PPP Subscriber Management 692
Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover 693
Verifying the Configuration of Access and Access-Internal Routes for Subscriber Management 694

Chapter 36 **Dynamic Profile Examples 695**

Example: IGMP Dynamic Profile 695
Example: Firewall Dynamic Profile 696
Example: Minimum MLPPP Dynamic Profile 697
Example: Minimum PPPoE Dynamic Profile 697

 Copyright © 2013, Juniper Networks, Inc.

Part 11 Dynamic VLANs

Chapter 37 Dynamic VLAN Overview 701

Dynamic 802.1Q VLAN Overview 701
Static VLAN Configuration 701
Dynamic VLAN Configuration 701
Agent Circuit Identifier-Based Dynamic VLANs Overview 702
VLAN Architectures and Subscriber Identification 702
ACI-Based Dynamic VLANs and Agent Circuit Identifier Interface Sets 703
Agent Circuit Identifier-Based Dynamic VLANs Components Overview 704
ACI-Based Dynamic Subscriber Interface 704
Dynamic ACI Interface Set 704
Dynamic or Static Underlying VLAN Interface 705
Static Physical Interface 705
Agent Circuit Identifier-Based Dynamic VLANs Bandwidth Management Overview 706
CoS Shaping Rate Adjustment 706
CoS Overhead Accounting Adjustment 707
Dynamic Profiles and Adjustment of CoS Shaping Rate and Overhead Accounting 707
Guidelines for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting 708
Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ACI Interface Sets 709
Ethernet OAM Support for Service VLANs Overview 709
Ethernet OAM Support for Service VLANs Terms and Acronyms 710
Components of Ethernet OAM Support for Service VLANs 710
How Ethernet OAM Support for Service VLANs Works 711
Restrictions for Using Ethernet OAM Support for Service VLANs 712

Chapter 38 Configuring Dynamic VLANs 713

Configuring VLAN Dynamic Profiles 714
Configuring a VLAN Dynamic Profile for Creating Single-Tag VLANs Using Standard TPID Values 714
Configuring a VLAN Dynamic Profile for Creating Single-Tag VLANs Using Any TPID Values 716
Configuring a Stacked VLAN Dynamic Profile 717
Configuring a VLAN Dynamic Profile That Associates VLAN Interfaces with Separate Routing Instances 718
Configuring VLAN Interfaces to Use Dynamic Profiles 720
Associating a Single-Tag VLAN Dynamic Profile with an Interface 720
Associating a Stacked VLAN Dynamic Profile with an Interface 721
Configuring Which VLAN Ethernet Packet Types Dynamic Profiles Can Accept 721
Configuring the VLAN Ethernet Packet Type for Single-Tag VLAN Dynamic Profiles 721
Configuring the VLAN Ethernet Packet Type for Stacked VLAN Dynamic Profiles 722
Configuring an Authentication Password for VLAN or Stacked VLAN Ranges . . 723

Copyright © 2013, Juniper Networks, Inc.

Configuring VLAN Ranges for Use with Dynamic Profiles 724
Configuring Single-Level VLAN Ranges for Use with VLAN Dynamic Profiles 724
Configuring Stacked VLAN Ranges for Use with Stacked VLAN Dynamic Profiles 725
Configuring Dynamic Mixed VLAN Ranges 726
Configuring VLAN Dynamic Profile Override 727
Configuring Dynamic Authentication for VLAN Interfaces 728
Configuring VLAN Interface Username Information for AAA Authentication . . . 729
Option 82 Suboptions in Authentication Usernames for Autosense VLANs 730
Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs 731
Automatically Removing VLANs with No Subscribers 731
Configuring Dynamic VLANs Based on Agent Circuit Identifier Information 733
Defining Agent Circuit Identifier Interface Sets 735
Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information 737
Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information 739
Configuring Dynamic VLAN Subscriber Interfaces Based on Agent Circuit Identifier Information 740
Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs 742
Configuring Ethernet OAM Support for Service VLANs with Double-Tagged Customer VLANs 743
Verifying and Managing Dynamic VLAN Configuration 747
Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration 747
Clearing Agent Circuit Identifier Interface Sets 748

Chapter 39 Dynamic VLAN Examples 751

Example: Configuring a VLAN Dynamic Profile for VLANs with a TPID of 0x8100 751
Example: Configuring a VLAN Dynamic Profile for VLANs with Any TPID Value and Enabling Demux Interfaces over the VLAN Interface 751
Example: Configuring a Stacked VLAN Dynamic Profile 752
Example: Dynamic VLAN Interface Configuration 752
Example: Dynamic Stacked VLAN Interface Configuration 753
Example: Dynamic Flexible VLAN Interface Configuration 753
Example: Configuring a Flexible VLAN Interface for Use with a Nonstandard Ethertype 754

Part 12 Subscriber Interfaces

Chapter 40 Subscriber Interface Overview 757

Subscriber Interface Overview 757
Statically Identifying Subscribers 757
Dynamically Identifying Subscribers 758
Static Subscriber Interfaces and VLAN Overview 758

 Copyright © 2013, Juniper Networks, Inc.

Subscriber Interfaces and Demultiplexing Overview 759
Interface Sets of Static Demux Interfaces 759
Dynamic Demultiplexing Interfaces 759
Guidelines for Configuring Demux Interfaces for Subscriber Access 760
IP Demux Interfaces over Static or Dynamic VLAN Demux Interfaces 760
MAC Address Validation for Subscriber Interfaces Overview 761
Supported Types of Subscriber Interfaces 762
Trusted Addresses 762
Types of MAC Address Validation 762

Chapter 41 **Configuring Subscriber Interfaces for Dynamic Profiles 765**

Configuring Static Subscriber Interfaces in Dynamic Profiles 765
Configuring a Subscriber Interface with a Static VLAN Interface 766
Configuring Static Subscriber Interfaces Using IP Demux Interfaces 766
Configuring Static Subscriber Interfaces Using VLAN Demux Interfaces . . . 767
Associating Dynamic Profiles with Statically Created Interfaces 768
Configuring a Subscriber Interface Using a Set of Static IP Demux Interfaces . . 769
Configuring a Subscriber Interface Using a Set of Static VLAN Demux Interfaces 770
Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles 771
Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles 772
Configuring MAC Address Validation for Subscriber Interfaces 773
Configuring MAC Address Validation for Static Subscriber Interfaces 774
Configuring MAC Address Validation for Dynamic Subscriber Interfaces . . . 775
Verifying Configuration and Status of Dynamic Subscribers and Associated Sessions, Services, and Firewall Filters 776

Chapter 42 **Subscriber Interface Examples 779**

Example: Configuring a Static Subscriber Interface on a Gigabit Ethernet VLAN Interface (Multiple Logical Units) 779
Example: Configuring a Static Subscriber Interface on a Gigabit Ethernet VLAN Interface 780
Example: Configuring a Static Subscriber Interface on a Gigabit Ethernet VLAN Interface (No Autonegotiation) 780
Example: Configuring a Static Subscriber Interface with a Loopback 780
Example: Configuring Dynamic Subscriber Interfaces on IP Demux Interfaces . . 781
Example: Configuring IPv6 Addressing for a Dynamic IP Demux Interface over Static VLANs 783
Example: Configuring IPv6 Addressing for a Dynamic IP Demux Interface over Dynamic VLANs 784
Example: Configuring a Dynamic IP Demux Interface with Dual Stacking 787
Example: Configuring IPv4 Static VLAN Demux Interfaces over a Gigabit Ethernet Underlying Interface with DHCP Local Server 791

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring IPv4 Dynamic VLAN Demux Interfaces over a Gigabit Ethernet Underlying Interface with DHCP Local Server 793
Example: Dynamic IP Demux Subscriber Interfaces over Dynamic VLAN Demux Interfaces 795
Example: Concurrent Configuration of Dynamic DHCP IP Demux and PPPoE Demux Interfaces over the Same VLAN Demux Interface 801
Example: Configuring CoS on Static LSQ MLPPP Bundle Interfaces 811

Chapter 43 **Subscriber Interfaces over Aggregated Ethernet Overview 815**

Static and Dynamic VLAN Subscriber Interfaces over Aggregated Ethernet Overview 815
Guidelines for Configuring an Aggregated Ethernet Logical Interface to Support a Static or Dynamic VLAN Subscriber Interface 815
Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview 816
Options for Aggregated Ethernet Logical Interfaces That Support Demux Subscriber Interfaces 816
Hardware Requirements with Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet 817
Features Supported with Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet 817
Distribution of Demux Subscribers in an Aggregated Ethernet Interface 818
Distribution Models 818
Sample Targeted Distribution Topology 819
Redundancy and Redistribution Mechanisms 819
Considerations and Best Practices 820

Chapter 44 **Configuring Subscriber Interfaces over Aggregated Ethernet 823**

Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet 824
Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet 825
Configuring a Static or Dynamic VLAN Demux Subscriber Interface over Aggregated Ethernet 826
Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces 828
Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface 829
Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface 829
Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface 830
Configuring Manual Rebalancing of Subscribers on an Aggregated Ethernet Interface 830
Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface 830
Configuring the PPPoE Family for an Underlying Interface 831

 Copyright © 2013, Juniper Networks, Inc.

Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces 833
Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface 833

Chapter 45 **Subscriber Interfaces over Aggregated Ethernet Examples 835**

Example: Configuring a Static Subscriber Interface on a VLAN Interface over Aggregated Ethernet 835
Example: Configuring a Static Subscriber Interface on an IP Demux Interface over Aggregated Ethernet 838
Example: Configuring a Static Subscriber Interface on a VLAN Interface over Aggregated Ethernet 840
Example: Configuring IPv4 Static VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server 843
Example: Separating Targeted Multicast Traffic for Demux Subscribers on Aggregated Ethernet Interfaces 846
Example: Configuring IPv4 Dynamic VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server 856
Example: Configuring IPv6 Dynamic VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server 859
Example: Configuring IPv4 Dynamic Stacked VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server 862
Example: Configuring a Static PPPoE Subscriber Interface on a Static Underlying VLAN Demux Interface over Aggregated Ethernet 865
Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Underlying VLAN Demux Interface over Aggregated Ethernet 870
Example: Configuring a Dynamic PPPoE Subscriber Interface on a Dynamic Underlying VLAN Demux Interface over Aggregated Ethernet 876

Chapter 46 **Dynamic PPPoE Subscriber Interfaces Overview 883**

Subscriber Interfaces and PPPoE Overview 883
Benefits of Using Dynamic PPPoE Subscriber Interfaces 884
Supported Platforms for Dynamic PPPoE Subscriber Interfaces 885
Sequence of Operations for PPPoE Subscriber Access 885
Sequence When a PPPoE Subscriber Logs In 885
Sequence When a PPPoE Subscriber Logs Out 886
Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview 887
PPPoE Dynamic Profile Configuration 887
PPPoE Underlying Interface Configuration 888
Address Assignment for Dynamic PPPoE Subscriber Interfaces 888
Guidelines for Configuring Dynamic PPPoE Subscriber Interfaces 889
PPPoE Maximum Session Limit Overview 890
Per-Interface Configuration for PPPoE Maximum Session Limit Using the CLI 890
Per-Subscriber Configuration for PPPoE Maximum Session Limit Using RADIUS 890
Override of PPPoE Maximum Session Limit from RADIUS 891
Guidelines for Using PPPoE Maximum Session Limit from RADIUS 892

Copyright © 2013, Juniper Networks, Inc.

PPPoE Subscriber Session Lockout Overview 893
Benefits of Using PPPoE Subscriber Session Lockout 893
Supported Platforms and Underlying Interfaces for PPPoE Subscriber Session Lockout 894
How PPPoE Subscriber Session Lockout Works 894
PPPoE Subscriber Session Lockout Period 895
PPPoE Subscriber Session Lockout and Duplicate Protection 895
PPPoE Subscriber Session Lockout and Automatic Removal of Dynamic Subscriber VLANs 896
Understanding the Lockout Period for PPPoE Subscriber Session Lockout 896
Duration of PPPoE Subscriber Session Lockout Period 896
How the Router Determines the PPPoE Subscriber Session Lockout Period 897

Chapter 47 Configuring Dynamic PPPoE Subscriber Interfaces 899

Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles . . . 899
Configuring a Basic PPPoE Dynamic Profile 900
Configuring a PPPoE Dynamic Profile with Additional Options 903
Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces 905
Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface 908
Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation 910
Configuring Lockout of PPPoE Subscriber Sessions 912
Clearing Lockout of PPPoE Subscriber Sessions 913
Verifying and Managing Dynamic PPPoE Configuration 914

Chapter 48 Dynamic PPPoE Subscriber Interfaces Examples 917

Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Gigabit Ethernet VLAN Interface 917
Example: Configuring a PPPoE Service Name Table for Dynamic Subscriber Interface Creation 919
Evaluation Order for Matching Client Information in PPPoE Service Name Tables 922

Chapter 49 MPLS Pseudowire Subscriber Interfaces Overview 923

Pseudowire Subscriber Logical Interfaces Overview 923
Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview 925
CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces 925
CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces 927
Three-Level Scheduling Hierarchy: Pseudowire Logical Interfaces over a Transport Logical Interface 928
Three-Level Scheduling Hierarchy : Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set 928
Three-Level Scheduling Hierarchy Combined Deployment Scenario 929
CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces 930

 Copyright © 2013, Juniper Networks, Inc.

Chapter 50 Configuring MPLS Pseudowire Subscriber Interfaces 933

Configuring a Pseudowire Subscriber Logical Interface 933
Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router 935
Configuring a Pseudowire Subscriber Logical Interface Device 935
Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface 937
Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces 937
Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces 938
Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface 939
Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces 940
Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Transport Logical Interface) 942
Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Pseudowire Interface Set) 944

Part 13 Class of Service for Subscriber Access

Chapter 51 Dynamic CoS for Subscriber Access Overview 949

CoS for Subscriber Access Overview 949
Guidelines for Configuring Dynamic CoS for Subscriber Access 950
Hardware Requirements for Dynamic CoS 950
Configuration Guidelines for Dynamic Scheduling and Queuing 953
Configuration Guidelines for Dynamic Classifiers and Rewrite Rules 953
Configuration Guidelines for Dynamic CoS on Specific Interface Types . . . 956
Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers 956
Two-Level Hierarchical Scheduling 957
Three-Level Hierarchical Scheduling 958
Interface Hierarchy Versus CoS Hierarchy 960

Chapter 52 Configuration Summary of Dynamic CoS for Subscriber Access 963

Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access 963
Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access 965
Configuring Per-Unit Scheduling in a Dynamic Profile for Subscriber Access . . . 967

Chapter 53 Configuring Dynamic Shaping and Scheduling for Subscriber Access . . . 971

Configuring Traffic Scheduling and Shaping for Subscriber Access 971
Configuring Static Traffic Shaping and Scheduling Parameters in a Dynamic Profile 971
Configuring Dynamic Traffic Shaping and Scheduling Parameters in a Dynamic Profile 972
Configuring Schedulers in a Dynamic Profile for Subscriber Access 973
Configuring Static Schedulers in a Dynamic Profile 974
Configuring Dynamic Schedulers with Variables in a Dynamic Profile 975
Configuring a Combination of Static and Dynamic Scheduler Parameters in a Scheduler Definition 976
Applying CoS Parameters to a Subscriber Interface in a Dynamic Profile 979
Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile 979
Applying Minimal Shaping and Scheduling to Remaining Subscriber Traffic 980
Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile 980
Applying a Classifier to a Subscriber Interface in a Dynamic Profile 981
Configuring Scheduler and Scheduler Map Sharing 982
Applying CoS Attributes to VLANs Using Agent-Circuit-Identifiers 984
Verifying the Scheduling and Shaping Configuration for Subscriber Access . . . 986

Chapter 54 RADIUS and Dynamic CoS Overview 989

Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS 989
Dynamic Configuration of Initial CoS in Access Profiles 989
Predefined Variables for Dynamic Configuration of Initial Traffic Shaping 990
Predefined Variables for Dynamic Configuration of Initial Scheduling and Queuing 991
Changing CoS Services Overview 993
Types of CoS Variables Used in a Service Profile 993
Static and Dynamic CoS Configurations 993
Scenarios for Static and Dynamic Configuration of CoS Parameters 994
CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview 996
Supported Network Configurations 996
Traffic-Control Profiles in Subscriber Interface Dynamic Profiles 996
CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets and Member Subscriber Sessions 997
Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions 998

Chapter 55 Configuring RADIUS for Dynamic CoS 999

Configuring Initial CoS Parameters Dynamically Obtained from RADIUS 999
Configuring User-Defined CoS Variables in a Dynamic Service Profile 1000

 Copyright © 2013, Juniper Networks, Inc.

Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions . . . 1003
CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets . . . 1005

Chapter 56 **Interface Solutions for Dynamic CoS Overview . . . 1013**

CoS for Aggregated Ethernet Subscriber Interfaces Overview . . . 1013
Guidelines for Configuring CoS for Aggregated Ethernet Subscribers . . . 1013
CoS for PPPoE Subscriber Interfaces Overview . . . 1014
CoS for L2TP LAC Subscriber Interfaces Overview . . . 1016
Traffic from LAC to LNS . . . 1016
LAC Tunnels: Traffic from LNS to LAC . . . 1017
CoS for L2TP LNS Inline Services Overview . . . 1018
Guidelines for Applying CoS to the LNS . . . 1018
Hardware Requirements for Inline Services on the LNS . . . 1018
CoS for Interface Sets of Subscribers Overview . . . 1019
Guidelines for Configuring Dynamic Interface Sets in a Subscriber Access Network . . . 1019

Chapter 57 **Configuring Interface Solutions for Dynamic CoS . . . 1023**

Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links . . . 1023
Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface . . . 1024
Configuring Dynamic CoS for an L2TP LAC Tunnel . . . 1025
Configuring Dynamic CoS for an L2TP LNS Inline Service . . . 1026
Configuring an Interface Set of Subscribers in a Dynamic Profile . . . 1028

Chapter 58 **Dynamic CoS for Subscriber Access Examples . . . 1031**

Example: Configuring Static Hierarchical Scheduling and Queuing for Subscriber Access . . . 1031
Example: Configuring Dynamic Hierarchical Scheduling and Queuing for Subscriber Access . . . 1033
Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS . . . 1040
Example: Providing Unique Rate Configurations for Schedulers in a Dynamic Profile . . . 1042
Example: Configuring Aggregate Scheduling of Queues for Residential Subscribers on Static IP Demux Interfaces . . . 1043
Example: Configuring Hierarchical Scheduling and Queuing for a Static PPPoE Subscriber Interface . . . 1045
Example: Configuring Hierarchical Scheduling and Queuing for an Underlying Static PPPoE Subscriber Interface . . . 1047
Example: Configuring Hierarchical Scheduling and Queuing for an Interface Set of Static PPPoE Subscriber Interfaces . . . 1050
Example: Configuring a Dynamic Interface Set of VLAN Subscribers . . . 1052
Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile . . . 1065

Copyright © 2013, Juniper Networks, Inc.

Chapter 59 Bandwidth Management for Dynamic CoS Overview 1069
Excess Bandwidth Distribution on MIC and MPC Interfaces Overview 1069
Traffic Burst Management on MIC and MPC Interfaces Overview 1070
Guidelines for Configuring the Burst Size 1070
How the System Calculates the Burst Size 1071
Bandwidth Management for Downstream Traffic in Edge Networks Overview 1072
Guidelines for Configuring the Shaping Mode 1072
Guidelines for Configuring Byte Adjustments 1073
Relationship with Other CoS Features 1073
Setting the Shaping-Rate and Overhead-Accounting Class-of-Service Attributes Based on Access Line Information 1073
Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags 1074
CLI Interaction with PPPoE Vendor-Specific Tags 1074
RADIUS Interaction with PPPoE Vendor-Specific Tags 1075
ANCP Interaction with PPPoE Vendor-Specific Tags 1075
Multicast QoS Adjustment Interaction with PPPoE Vendor-Specific Tags 1075
Shaping Rate Restrictions 1075
Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview 1076
Queue Scaling for MIC and MPC Combinations 1076
Distribution of Queues on 30-Gigabit Ethernet Queuing MPCs 1077
Distribution of Queues on 60-Gigabit Ethernet MPCs 1078
Determining Maximum Egress Queues and Subscriber Interfaces per Port 1079
Managing Remaining Queues 1079
Hierarchical CoS Shaping-Rate Adjustments Overview 1080
Types of Shaping-Rate Adjustments 1080
Levels of Shaping-Rate Adjustments 1080
Shaping Rate Adjustments for Subscriber Local Loops Overview 1082
Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops 1083
CoS Adjustment Control Profiles Overview 1084
Effective Shaping Rate 1085
Applications and Associated Algorithms in Adjustment Control Profiles . . 1085
Chapter 60 Configuring Bandwidth Management Parameters for Dynamic CoS . . . 1089
Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces 1090
Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates 1092

 Copyright © 2013, Juniper Networks, Inc.

Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces 1093
Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces 1093
Configuring the Maximum Number of Queues for MIC and MPC Interfaces 1093
Configuring Remaining Common Queues on MIC and MPC Interfaces . . . 1094
Verifying the Number of Dedicated Queues Configured on MIC and MPC Interfaces 1095
Providing Accurate Scheduling for a Demux Subscriber Interface of Aggregated Ethernet Links 1095
Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers 1096
Overview 1097
Configuring a Static Minimum Adjusted Shaping Rate on Scheduler Nodes 1097
Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes 1097
Configuring Shaping-Rate Adjustments on Queues 1097
Overview 1098
Configuring a Static Shaping-Rate Adjustment for Queues 1098
Configuring a Dynamic Shaping-Rate Adjustment for Queues 1099
Enabling Shaping-Rate Adjustments for Subscriber Local Loops 1099
Configuring Static Logical Interface Sets to Serve as CoS Hierarchical Scheduler Nodes for Subscriber Loops 1100
Configuring the Logical Interfaces That Compose the Static Logical Interface Sets 1101
Configuring Hierarchical CoS on the Static Logical Interface Sets That Serve as Hierarchical Scheduler Nodes for Subscriber Local Loops 1101
Configuring ANCP Functionality That Supports and Drives Shaping-Rate Adjustments for Subscriber Local Loops 1103
Disabling Shaping-Rate Adjustments for Subscriber Local Loops 1104
Disabling Hierarchical Bandwidth Adjustment for Subscriber Interfaces with Reverse-OIF Mapping 1105
Verifying the Configuration of Shaping-Rate Adjustments for Subscriber Local Loops 1105
Verifying the Configuration of ANCP for Shaping-Rate Adjustments 1106
Configuring CoS Adjustment Control Profiles 1106
Verifying the CoS Adjustment Control Profile Configuration 1107
Reporting the Effective Shaping Rate for Subscribers 1108
Verifying the Effective Shaping Rate Reporting Configuration 1108

Chapter 61 Bandwidth Management for Dynamic CoS Examples 1111

Example: Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates 1111
Managing Traffic with Different Encapsulations 1111
Managing Downstream Cell-Based Traffic 1113
Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops 1114

Part 14 Subscriber Service Activation

Chapter 62 Subscriber Service Activation Overview 1121

CLI-Activated Subscriber Services 1121
Activating and Deactivating Subscriber Services Locally with the CLI 1122
Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers 1125
Using the CLI to Globally Modify a Traffic-Control Profile Currently Applied to Multiple Subscribers 1126
Using the CLI to Modify a Traffic-Control Profile for a Specific Current Subscriber 1126
Default Subscriber Service Overview 1127
Configuring a Default Subscriber Service 1127

Part 15 Dynamic Firewall Filters, Service Sets and HTTP Redirect for Subscriber Access

Chapter 63 Dynamic Firewall Filters and Service Sets Overview 1131

Dynamic Firewall Filters Overview 1132
Classic Filters Overview 1133
Classic Filter Types 1133
Classic Filter Components 1133
Classic Filter Processing 1134
Guidelines for Creating and Applying Classic Filters for Subscriber Interfaces 1135
Basic Classic Filter Syntax 1135
Ascend-Data-Filter Policies for Subscriber Management Overview 1136
Filter Naming Conventions 1137
Use of Multiple Sessions with Ascend-Data-Filters on an Interface 1137
Optional ADF Filter Requirement for Some Subscribers 1138
Ascend-Data-Filter Attribute Fields 1138
Firewall Filters and Enhanced Network Services Mode Overview 1141
Hierarchical Policer Overview 1142
Hierarchical Policer as Filter Action 1143
Interface-Shared Filters Overview 1144
Enhanced Policer Statistics Overview 1144
Fast Update Filters Overview 1146
Fast Update Filter Components 1147
Fast Update Filter Processing 1147
Fast Update Filter Names 1148
Guidelines for Creating and Applying Fast Update Filters 1148
Basic Fast Update Filter Syntax 1149
Match Conditions and Actions in Fast Update Filters 1150
Match Conditions 1150
Actions 1151
Adding Terms Only Once 1151
Dynamic Service Sets Overview 1151

 Copyright © 2013, Juniper Networks, Inc.

Chapter 64 Parameterized Filters 1153
Parameterized Filters Overview 1153
Basic Parameterized Filter Syntax 1154
Unique Identifiers for Firewall Variables in Dynamic Profiles 1154
Sample Dynamic-Profile Configuration for Parameterized Filters 1156
Dynamic Profile Configuration and UID Substitution Comparison for Parameterized Filters 1158
Example: Dynamic-Profile Parsing 1162
Parameterized Filters Configuration Considerations 1163
Subscriber IP Address 1163
Interaction with Static Configuration 1163
Interface-Specific 1164
Service Session Support 1164
Filter Naming Conventions 1164
Guidelines for Creating and Applying Parameterized Filters for Subscriber Interfaces 1164
Parameterized Filter Processing Overview 1165
IPv4 Parameterized Filter Match Conditions 1166
IPv6 Parameterized Filter Match Conditions 1167
Parameterized Filter Actions and Modifiers 1168
Parameterized Filter Policer Actions 1168

Chapter 65 Configuring Filters for Dynamic Profiles 1169
Dynamically Attaching Statically Created Filters for a Specific Interface Family Type 1169
Dynamically Attaching Statically Created Filters for Any Interface Type 1170
Dynamically Attaching Filters Using RADIUS Variables 1172
Defining Dynamic Filter Processing Order 1173
Configuring Firewall Filter Bypass 1174
Configuring Service Packet Counting 1175
Configuring a Filter for Use with Enhanced Network Services Mode 1176
Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions 1177
Verifying and Managing Dynamic Ascend-Data-Filter Policy Configuration 1179
Configuring Unicast RPF and Fail Filters in Dynamic Profiles for Subscriber Interfaces 1179
Configuring Unicast RPF in Dynamic Profiles for Subscriber Interfaces 1180
Configuring a Fail Filter for Unicast RPF in Dynamic Profiles for Subscriber Interfaces 1180

Chapter 66 Configuring Fast Update Filters 1183
Configuring Fast Update Filters 1183
Configuring the Match Order for Fast Update Filters 1184
Configuring Terms for Fast Update Filters 1185
Fast Update Filter Match Conditions 1186
Fast Update Filter Actions and Action Modifiers 1187
Configuring Filters to Permit Expected Traffic 1187
Avoiding Conflicts When Terms Match 1188
How the Router Evaluates Terms in a Filter 1189
Using Implied Wildcards 1190

Copyright © 2013, Juniper Networks, Inc.

Conflict Caused by Overlapping Ranges 1191
Associating Fast Update Filters with Interfaces in a Dynamic Profile 1193
Verifying and Managing Firewall Filter Configuration 1194

Chapter 67 **Configuring Service Sets in Dynamic Profiles 1197**

Associating Service Sets with Interfaces in a Dynamic Profile 1197
Verifying and Managing Service Sets Information 1198

Chapter 68 **Firewall Filter Examples 1199**

Examples: Configuring Static Filters 1199
Example: Configuring Dynamic Ascend-Data-Filter Support for Subscriber Access 1202
Example: Configuring Static Ascend-Data-Filter Support for Subscriber Access 1205
Example: Configuring Fast Update Filters for Subscriber Access 1209
Example: Bypassing Firewall Filters 1210
Example: Configuring Hierarchical Policers as Filter Actions 1214
Example: Interface-Shared Filter Configuration 1217
Example: Configuring Unicast RPF in a Dynamic Profile on MX Series Routers . . 1219

Chapter 69 **Redirecting HTTP Requests Overview 1225**

Redirecting HTTP Requests Overview 1225
Remote HTTP Redirect Server Operation Flow 1226
Local HTTP Redirect Server Operation Flow 1227

Chapter 70 **Configuring HTTP Redirect 1229**

Configuring HTTP Redirect Services 1229
Verifying HTTP Redirect Requests 1232

Chapter 71 **HTTP Redirect Examples 1233**

Example: Walled Garden as a Service Filter 1233
Example: Walled Garden as an HTTP Service Rule 1234
Example: HTTP Service Within a Service Set 1235
Example: HTTP Service Attached to a Static Interface 1235
Example: HTTP Service Attached to a Dynamic Interface 1237
Example: Configuring Redundant Multiservice 1239
Example: Configuring Destination Address Rewrite for HTTP Redirect 1240

Part 16 Subscriber Secure Policy Traffic Mirroring

Chapter 72 **Subscriber Secure Policy Traffic Mirroring Overview 1245**

Subscriber Secure Policy Overview 1245
Subscriber Secure Policy for Subscribers on VLANs 1245
Traffic Filtering For DTCP-Initiated Subscriber Secure Policy Mirrored Traffic 1246
Mirroring-Related Event Reporting 1246
Subscriber Secure Policy Licensing Requirements 1246

Chapter 73 **Overview of Subscriber Secure Policy Using RADIUS 1247**

RADIUS-Initiated Subscriber Secure Policy Overview 1247
Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS 1248
RADIUS-Initiated Traffic Mirroring Interfaces 1250

 Copyright © 2013, Juniper Networks, Inc.

RADIUS-Initiated Traffic Mirroring Process at Subscriber Login 1251
RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers 1252
Subscriber Secure Policy Support for IPv4 Multicast Traffic 1253
Triggering the Mirroring of IPv4 Multicast Traffic 1253
RADIUS Attributes Used for Subscriber Secure Policy 1254
Triggering Subscriber Secure Policy for Subscribers on Dynamic Authenticated VLANs 1255
Using the Packet Header to Track Subscribers on the Mediation Device 1255
Format of the Mirror Header Values Used to Track Subscribers and Subscriber Sessions 1258
4-Byte Format 1258
8-Byte Format 1258
Subscriber Secure Policy and L2TP LAC Subscribers 1259
Subscriber Secure Policy and L2TP LNS Subscribers 1259

Chapter 74 Configuring RADIUS-Initiated Subscriber Secure Policy Traffic Mirroring 1261

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview . . . 1261
Guidelines for Configuring Subscriber Secure Policy Mirroring 1262
Configuring Tunnel Interfaces for Subscriber Secure Policy Mirroring 1263
Configuring Support for Subscriber Secure Policy Mirroring 1264
Configuring RADIUS Server Support for Subscriber Secure Policy Mirroring . . . 1265
Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic 1266
Terminating RADIUS-Initiated Subscriber Traffic Mirroring 1266

Chapter 75 Overview of Subscriber Secure Policy Using DTCP 1267

DTCP-Initiated Subscriber Secure Policy Overview 1267
Subscriber Secure Policy Traffic Mirroring Architecture Using DTCP 1267
DTCP-Initiated Traffic Mirroring Interfaces 1269
DTCP-Initiated Traffic Mirroring Process 1271
Subscriber Secure Policy Support for IPv4 Multicast Traffic 1272
Triggering the Mirroring of IPv4 Multicast Traffic 1272
Subscriber Secure Policy and L2TP LNS Subscribers 1272
DTCP Messages Used for Subscriber Secure Policy 1273
DTCP Traffic Mirroring Triggers 1273
Triggering Subscriber Secure Policy for Subscribers on Dynamic Authenticated VLANs 1276
Order in Which Trigger Attributes Are Processed 1276
Packet Header for Mirrored Traffic Sent to Mediation Device 1278

Chapter 76 Configuring DTCP-Initiated Subscriber Secure Policy Mirroring 1281

Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview 1281
Guidelines for Configuring Subscriber Secure Policy Mirroring 1282
Configuring Tunnel Interfaces for Subscriber Secure Policy Mirroring 1283
Configuring Support for Subscriber Secure Policy Mirroring 1284
Example: Configuring Traffic That Is Mirrored Using DTCP-Initiated Subscriber Secure Policy 1285
Configuring the Mediation Device as a User on the Router 1287
Configuring a DTCP-over-SSH Connection to the Mediation Device 1288
Configuring the Mediation Device to Provision Traffic Mirroring 1288

Copyright © 2013, Juniper Networks, Inc.

Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic 1289
Terminating DTCP-Initiated Subscriber Traffic Mirroring Sessions 1289

Chapter 77 **DTCP Messages Used for DTCP-Initiated Subscriber Secure Policy Mirroring 1291**

ADD DTCP 1292
DELETE DTCP 1295
DISABLE DTCP 1297
ENABLE DTCP 1298
LIST DTCP 1299
Example: Using DTCP Messages to Trigger, Verify, and Disable Traffic Mirroring for Subscribers 1300
Creating DTCP ADD Messages to Trigger Traffic Mirroring 1301
Creating DTCP ENABLE Messages to Trigger Traffic Mirroring 1302
Creating DTCP DISABLE Messages to Trigger Traffic Mirroring 1302
Using LIST Messages to Verify That Subscriber Traffic Is Being Mirrored . . 1302
Using DELETE Messages to Remove Traffic Mirroring Triggers 1303
Using Disable Messages to Disable Traffic Mirroring Triggers 1304
Using Enable Messages to Enable Traffic Mirroring Triggers 1304
Verifying That Traffic Mirroring Was Stopped on the Subscriber Interfaces 1304

Chapter 78 **Reporting Intercept-Related Events for Subscriber Secure Policy 1305**

Intercept-Related Events Transmitted to the Mediation Device 1305
SNMP Traps for Subscriber Secure Policy LAES Compliance 1305
Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring 1307
Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring 1307

Part 17 Protocols for Subscriber Access

Chapter 79 **ANCP Overview 1311**

ANCP Topology Discovery and Traffic Reporting Overview 1311
Overview 1311
Topology Discovery 1312
Subscriber Services 1312
ANCP Interfaces and Access Loop Circuit Identifiers 1313
ANCP Neighbors 1313
Traffic Rate Reporting and Adjustment by ANCP 1314
Overview 1315
Traffic Rate Adjustment 1315
Recommended Traffic Shaping Rates 1316
ANCP Keepalives 1316
ANCP Interactions with AAA 1317
ANCP Operations in Different Network Configurations 1319
ANCP Network Using N:1 and 1:1 Configuration Models without Interface Sets 1322
Sequence of ANCP Events: Static VLAN or VLAN Demux Interfaces over Ethernet Without Interface Sets 1323
ANCP Network Using N:1 Configuration Model with Interface Sets 1324

 Copyright © 2013, Juniper Networks, Inc.

Sequence of ANCP Events: Static VLAN Interfaces over Ethernet with Interface Sets . . . 1325
ANCP Network Using 1:1 Configuration Model with Interface Sets . . . 1326
Sequence of ANCP Events: Static VLAN Demux Interfaces over Aggregated Ethernet with Interface Sets . . . 1327
ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes . . . 1327
Preservation of CoS Shaping Across ANCP Restarts . . . 1329

Chapter 80 **Configuring ANCP . . . 1331**

Configuring ANCP . . . 1332
Configuring ANCP Neighbors . . . 1333
Associating an Access Node with Subscribers for ANCP Operations . . . 1334
Specifying the Interval Between ANCP Adjacency Messages . . . 1335
Specifying the Maximum Number of Discovery Table Entries . . . 1335
Configuring ANCP for Backward Compatibility . . . 1336
Specifying How Long Processes Wait for ANCP Restart to Complete . . . 1337
Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates . . . 1337
Configuring the ANCP Agent to Report Traffic Rates to CoS . . . 1338
Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces . . . 1339
Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages . . 1341
Triggering ANCP OAM . . . 1342
Verifying and Monitoring ANCP Neighbors . . . 1343
Clearing ANCP Neighbors . . . 1343
Verifying and Monitoring ANCP Subscribers . . . 1344
Clearing ANCP Subscribers . . . 1344
Verifying and Monitoring CoS for ANCP Subscribers . . . 1345
Tracing ANCP Operations for Subscriber Access . . . 1345
Configuring the ANCP Trace Log Filename . . . 1346
Configuring the Number and Size of ANCP Log Files . . . 1346
Configuring Access to the ANCP Log File . . . 1347
Configuring a Regular Expression for ANCP Messages to Be Logged . . . 1347
Configuring the ANCP Tracing Flags . . . 1348
Configuring the Severity Level to Filter Which ANCP Messages Are Logged . . . 1348
Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet . . . 1349

Chapter 81 **Dynamic IGMP Configuration Overview . . . 1371**

Dynamic IGMP Configuration Overview . . . 1371

Chapter 82 **Dynamic MLD Configuration Overview . . . 1373**

Dynamic MLD Configuration Overview . . . 1373

Chapter 83 **Dynamic Router Advertisement Overview . . . 1375**

Dynamic Router Advertisement Configuration Overview . . . 1375

Part 18 Subscriber Access Examples

Chapter 84 **Service Profile Examples . . . 1379**

Example: Configuring a Tiered Service Profile for Subscriber Access . . . 1379

Part 19 Complete Configuration Statement Hierarchy and Summary of Statements for Subscriber Access

Chapter 85 Subscriber Access Statement Hierarchy 1385

[edit access address-assignment] Hierarchy Level 1385
[edit access domain] Hierarchy Level 1386
[edit access gx-plus] Hierarchy Level 1387
[edit access profile] Hierarchy Level 1387
[edit access tunnel-profile] Hierarchy Level 1389
[edit access tunnel-switch-profile] Hierarchy Level 1390
[edit diameter] Hierarchy Level 1390
[edit dynamic-profiles] Hierarchy Level 1391
[edit forwarding-options dhcp-relay] Hierarchy Level 1398
[edit interfaces radius-options] Hierarchy Level 1402
[edit jsrc] Hierarchy Level 1402
[edit protocols ancp] Hierarchy Level 1403
[edit services captive-portal-content-delivery] Hierarchy Level 1403
[edit services l2tp] Hierarchy Level 1404
[edit services mobile-ip] Hierarchy Level 1405
[edit services radius-flow-tap] Hierarchy Level 1406
[edit system services dhcp-local-server] Hierarchy Level 1407
[edit system services packet-triggered-subscribers] Hierarchy Level 1411
[edit system services static-subscribers] Hierarchy Level 1411
[edit system services subscriber-management] Hierarchy Level 1412

Chapter 86 Subscriber Access Configuration Statements 1413

aaa-access-profile (L2TP LNS) 1413
aaa-logical-system (Domain Map) 1414
aaa-routing-instance (Domain Map) 1415
abated-utilization (Address-Assignment Pools) 1415
abated-utilization-v6 (Address-Assignment Pools) 1416
access (Dynamic Access Routes) 1417
access-concentrator 1418
access-identifier 1419
access-internal (Dynamic Access-Internal Routes) 1420
access-loop-id-local 1421
access-profile (Domain Map) 1421
access-profile (Routing Instances) 1422
access-profile (Static Subscribers) 1423
access-profile-name (Duplicate Accounting) 1424
access-type 1425
accounting (Access Profile) 1426
accounting (Dynamic IGMP Interface) 1427
accounting (Dynamic MLD Interface) 1427
accounting-backup-options (Access Profile) 1428
accounting-order (Service Accounting) 1428
accounting-port 1429
accounting-server 1429
accounting-session-id-format 1430
accounting-stop-on-access-deny 1430

Copyright © 2013, Juniper Networks, Inc.

accounting-stop-on-failure . . . 1431
action . . . 1432
active-server-group . . . 1433
address . . . 1434
address (Diameter Peer) . . . 1434
address (Diameter Transport) . . . 1435
address (Tunnel Profile Remote Gateway) . . . 1435
address (Tunnel Profile Source Gateway) . . . 1436
address-assignment (Address-Assignment Pools) . . . 1437
address-change-immediate-update . . . 1438
address-pool (Domain Map) . . . 1438
address-protection . . . 1439
adf (Dynamic Firewalls) . . . 1440
adjacency-timer . . . 1441
adjustment-control-profiles . . . 1442
adjust-minimum (Dynamic Shaping and Scheduling) . . . 1443
adjust-percent (Dynamic Schedulers) . . . 1444
advisory-options (Traffic Shaping) . . . 1445
agent-circuit-identifier (Dynamic VLAN interface Sets) . . . 1446
aggregate (Hierarchical Policer) . . . 1447
aggregate-clients (DHCP Local Server) . . . 1448
aggregate-clients (DHCP Relay Agent) . . . 1449
aggregate-clients (Static Subscribers) . . . 1451
algorithm . . . 1452
allow-snooped-clients . . . 1453
always-write-giaddr . . . 1454
always-write-option-82 . . . 1455
anchor-point (Pseudowire Subscriber Interfaces) . . . 1456
ancp . . . 1457
ancp (Adjustment Control Profiles) . . . 1458
application (Adjustment Control Profiles) . . . 1459
application . . . 1459
apply-groups (Subscriber Secure Policy) . . . 1460
apply-groups-except (Subscriber Secure Policy) . . . 1460
assignment-id-format (L2TP LAC) . . . 1461
attempts (DHCP Local Server) . . . 1462
attribute . . . 1463
attributes . . . 1463
authenticate . . . 1464
authentication (DHCP Local Server) . . . 1465
authentication (DHCP Relay Agent) . . . 1466
authentication (Static and Dynamic PPP) . . . 1467
authentication (Static Subscribers) . . . 1468
authentication-order . . . 1469
authentication-server . . . 1470
authorization-order . . . 1470
auto-configure . . . 1471
auto-configure (Dynamic VLAN Interface Sets) . . . 1472
autonomous (Dynamic Router Advertisement) . . . 1473

Copyright © 2013, Juniper Networks, Inc.

avp (L2TP Tunnel Switching) 1473
bandwidth (Inline Services) 1474
bandwidth-limit (Hierarchical Policer) 1475
bandwidth-limit (Policer) 1476
bandwidth-percent 1478
bearer-type (L2TP Tunnel Switching) 1480
bfd 1481
boot-file 1482
boot-server 1482
buffer-size (Dynamic Scheduling) 1483
burst-size-limit (Hierarchical Policer) 1484
burst-size-limit (Policer) 1485
bytes (Dynamic Traffic Shaping) 1488
calling-number (L2TP Tunnel Switching) 1489
calling-station-id-delimiter (Subscriber Management) 1490
calling-station-id-format (Subscriber Management) 1491
captive-portal-content-delivery 1492
captive-portal-content-delivery-rule 1493
captive-portal-content-delivery-rule-set 1493
cell-mode (Dynamic Traffic Shaping) 1494
challenge-length (Static and Dynamic PPP) 1495
chap 1496
chap (Dynamic PPP) 1497
chap (L2TP) 1497
circuit-id (Address-Assignment Pools) 1498
circuit-id (DHCP Relay Agent) 1499
circuit-type (DHCP Local Server) 1500
circuit-type (DHCP Relay Agent) 1501
cisco-nas-port-info (L2TP Tunnel Switching) 1502
class-of-service (Dynamic Profiles) 1503
classifiers (Dynamic CoS Application) 1503
clear-on-abort (DHCP Local Server) 1504
client 1505
client-accounting-algorithm 1506
client-authentication-algorithm 1507
client-discover-match (DHCP Local Server) 1508
client-discover-match (DHCP Relay Agent) 1509
client-id (DHCP Local Server) 1510
client-id (DHCP Relay Agent) 1511
client-idle-timeout 1511
client-session-timeout 1512
coa-dynamic-variable-validation 1512
coa-immediate-update 1513
coa-no-override service-class-attribute 1513
color-aware 1514
color-blind 1515
committed-burst-size 1516
committed-information-rate 1518
concurrent-data-sessions 1519

 Copyright © 2013, Juniper Networks, Inc.

connect-actively . . . 1520
current-hop-limit (Dynamic Router Advertisement) . . . 1520
database-replication (Subscriber Session Database) . . . 1521
default-action (DHCP Relay Agent Option) . . . 1522
default-lifetime (Dynamic Router Advertisement) . . . 1523
default-value . . . 1523
delay-buffer-rate (Dynamic Traffic Shaping) . . . 1524
delegated-pool (DHCP Local Server) . . . 1525
delimiter (DHCP Local Server) . . . 1526
delimiter (DHCP Relay Agent) . . . 1528
delimiter (Domain Map) . . . 1529
demux0 (Dynamic Interface) . . . 1530
demux-options (Dynamic Interface) . . . 1531
demux-source (Dynamic IP Demux Interface) . . . 1532
demux-source (Dynamic Underlying Interface) . . . 1533
destination (Diameter Network Element) . . . 1533
destination (L2TP Destination Lockout) . . . 1534
destination-address . . . 1534
destination-address (Subscriber Secure Policy) . . . 1535
destination-host . . . 1535
destination-host (Gx-Plus) . . . 1536
destination-host (PTSP) . . . 1536
destination-port (Subscriber Secure Policy) . . . 1537
destination-prefix-list . . . 1537
destination-realm (Gx-Plus) . . . 1538
destination-realm (JSRC) . . . 1538
destination-realm (PTSP) . . . 1539
destruct-timeout (L2TP) . . . 1539
detection-time . . . 1540
device-count (Pseudowire Subscriber Interfaces) . . . 1541
dhcp-attributes (Address-Assignment Pools) . . . 1542
dhcp-local-server . . . 1543
dhcp-relay . . . 1548
dhcpv6 (DHCP Local Server) . . . 1554
dhcpv6 (DHCP Relay Agent) . . . 1557
dial-options . . . 1560
dial-options (Dynamic Profiles) . . . 1561
diameter . . . 1562
diameter-instance (Gx-Plus) . . . 1563
diameter-instance (JSRC) . . . 1563
diameter-instance (PTSP) . . . 1564
disable . . . 1564
disable (Dynamic IGMP) . . . 1565
disable (Dynamic MLD) . . . 1565
disable-calling-number-avp (L2TP LAC) . . . 1566
disable-failover-protocol (L2TP LAC) . . . 1566
disable-relay . . . 1567
dns-server . . . 1567
domain (Domain Map) . . . 1568

Copyright © 2013, Juniper Networks, Inc.

domain-name (Address-Assignment Pools) 1569
domain-name (DHCP Local Server) 1570
domain-name (DHCP Relay Agent) 1572
domain-name (Static Subscribers) 1573
domain-name-server (Routing Instances and Access Profiles) 1574
domain-name-server-inet (Routing Instances and Access Profiles) 1575
domain-name-server-inet6 (Routing Instances and Access Profiles) 1576
downstream-rate (Traffic Shaping) 1577
drop (DHCP Relay Agent Option) 1578
drop-policy (Subscriber Secure Policy) 1579
drop-profile (Dynamic Schedulers) 1580
drop-profile-map (Dynamic Schedulers) 1581
dscp (Dynamic Classifiers) 1582
dscp (Dynamic Rewrite Rules) 1583
dscp (Subscriber Secure Policy) 1583
dscp-ipv6 (Dynamic Classifiers) 1584
dscp-ipv6 (Dynamic Rewrite Rules) 1584
duplicate-clients-on-interface (DHCP Local Server) 1585
duplicate-clients-on-interface (DHCP Relay Agent) 1585
duplicate-protection (Dynamic PPPoE) 1586
duplication (Access Profile) 1587
duplication-vrf (Duplicate Accounting) 1587
dynamic-class-of-service-options (Dynamic Traffic Shaping) 1588
dynamic-home-assignment 1589
dynamic-profile (DHCP Local Server) 1590
dynamic-profile (DHCP Relay Agent) 1591
dynamic-profile (Domain Map) 1592
dynamic-profile (Dynamic PPPoE) 1593
dynamic-profile (Dynamic VLAN Interface Sets) 1594
dynamic-profile (L2TP) 1594
dynamic-profile (PPP) 1595
dynamic-profile (PPPoE Service Name Tables) 1596
dynamic-profile (Static Subscribers) 1597
dynamic-profiles 1598
effective-shaping-rate 1605
enable 1605
enable-service 1606
enable-snmp-tunnel-statistics (L2TP) 1607
encapsulation (Dynamic Interfaces) 1608
encapsulation (Logical Interface) 1611
enforce-strict-scale-limit-license (Subscriber Management) 1615
enhanced-policer 1615
equals (DHCP Relay Agent Option) 1616
entity-type 1617
equals (Dynamic Profile) 1617
ethernet-port-type-virtual 1618
exceed-action 1618
excess-burst-size 1619
excess-priority (Dynamic Schedulers) 1620

 Copyright © 2013, Juniper Networks, Inc.

excess-rate (Dynamic Schedulers) 1621
excess-rate (Dynamic Traffic Shaping) 1622
excess-rate-high (Dynamic Traffic Shaping) 1623
excess-rate-low (Dynamic Traffic Shaping) 1624
exclude (RADIUS) 1625
exclude (Dynamic MLD Interface) 1629
external-authority 1629
fail-over-within-preference (L2TP LAC) 1630
failure-action 1631
family (Address-Assignment Pools) 1632
family (Dynamic Firewalls) 1633
family (Dynamic Demux Interface) 1634
family (Dynamic PPPoE) 1635
family (Dynamic Standard Interface) 1636
fast-update-filter (Dynamic Firewalls) 1638
filter (Configuring) 1639
filter (Dynamic Firewalls) 1640
filter (Dynamic Interface Unit) 1641
filter-specific 1642
firewall (Dynamic Firewalls) 1643
forward-only (DHCP Relay Agent Option) 1644
forward-snooped-clients (DHCP Local Server) 1645
forward-snooped-clients (DHCP Relay Agent) 1646
forwarding (Diameter Network Element) 1647
forwarding-class (Dynamic Scheduler Maps) 1647
forwarding-class (Subscriber Secure Policy) 1648
fpc (MX Series 3D Universal Edge Routers) 1649
frame-mode (Dynamic Traffic Shaping) 1651
from 1652
from (Subscriber Secure Policy) 1653
function (Diameter Network Element) 1653
function (Diameter Route) 1654
gateway-name (Tunnel Profile Remote Gateway) 1654
gateway-name (Tunnel Profile Source Gateway) 1655
generic 1655
global (Gx-Plus) 1656
grace-period 1656
gres-route-flush-delay (Subscriber Management) 1657
group (DHCP Local Server) 1658
group (DHCP Relay Agent) 1661
group (Dynamic IGMP Interface) 1664
group (Dynamic MLD Interface) 1665
group (Static Subscribers) 1666
group-count (Dynamic MLD Interface) 1667
group-increment (Dynamic MLD Interface) 1667
group-limit (Dynamic IGMP Interface) 1668
group-limit (Dynamic MLD Interface) 1669
group-policy (Dynamic IGMP Interface) 1669
group-policy (Dynamic MLD Interface) 1670

Copyright © 2013, Juniper Networks, Inc.

group-profile (Group Profile) 1671
guaranteed-rate (Dynamic Traffic Shaping) 1672
gx-plus (Gx-Plus) 1673
hardware-address 1673
hello-interval 1674
hierarchical-policer 1675
hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) 1677
high-utilization (Address-Assignment Pools) 1678
high-utilization-v6 (Address-Assignment Pools) 1679
holddown-interval 1680
home-agent (Mobile IP Dynamic Assignment) 1681
home-agent (Mobile IP Network Address Identifier) 1682
home-agent (Mobile IP Networks) 1683
home-agent-address 1684
host (Address-Assignment Pools) 1685
host (Diameter Origin) 1685
identification (Tunnel Profile) 1686
idle-timeout (Access) 1687
idle-timeout (L2TP) 1688
ieee-802.1 (Dynamic Classifiers) 1689
ieee-802.1 (Dynamic Rewrite Rules) 1690
ietf-mode 1690
if-exceeding (Policer) 1691
if-exceeding (Hierarchical Policer) 1692
igmp (Dynamic Profiles) 1693
ignore 1694
immediate-leave (Dynamic IGMP Interface) 1694
immediate-leave (Dynamic MLD Interface) 1695
immediate-update 1696
include-ipv6 (Gx-Plus) 1696
inet (Subscriber Secure Policy) 1697
inet6 (Subscriber Secure Policy) 1698
inet-precedence (Dynamic Classifiers) 1699
inet-precedence (Dynamic Rewrite Rules) 1699
inline-services (PIC level) 1700
inner-tag-protocol-id (Dynamic VLANs) 1700
inner-vlan-id (Dynamic VLANs) 1701
input (Dynamic Service Sets) 1702
input-vlan-map (Dynamic Interfaces) 1703
instance-role 1704
interface (DHCP Local Server) 1705
interface (DHCP Relay Agent) 1707
interface (Dynamic IGMP) 1709
interface (Dynamic Interface Sets) 1710
interface (Dynamic MLD) 1711
interface (Dynamic Routing Instances) 1712
interface (Dynamic Router Advertisement) 1713
interface (Dynamic Routing Options) 1714
interface (L2TP Service Interfaces) 1714

Copyright © 2013, Juniper Networks, Inc.

interface (Static Subscriber Group) 1715
interface (Static Subscriber Username) 1716
interface-client-limit (DHCP Local Server) 1717
interface-client-limit (DHCP Relay Agent) 1719
interface-delete (Subscriber Management or DHCP Client Management) 1720
interface-description-format 1721
interface-id 1722
interface-name (DHCP Local Server) 1723
interface-name (DHCP Relay Agent) 1724
interface-set (ANCP) 1725
interface-set (Dynamic CoS) 1726
interface-set (Dynamic VLAN Interface Sets Association) 1727
interface-set (Dynamic VLAN Interface Sets Definition) 1728
interface-shared 1729
interface-traceoptions (DHCP) 1730
interfaces (Subscriber Secure Policy) 1731
interfaces (ANCP) 1732
interfaces (Dynamic CoS Definition) 1733
interfaces (Static and Dynamic Subscribers) 1734
interface-specific (Dynamic Firewalls) 1738
ip-address 1738
ip-address-change-notify 1739
ip-address-first 1740
ip-reassembly (L2TP) 1741
ip-reassembly 1742
ip-reassembly-rules (Service Set) 1743
jsrc (JSRC) 1743
jsrc-partition 1744
juniper-dsl-attributes 1744
keepalive 1745
keepalives 1746
keepalives (Dynamic Profiles) 1747
key 1748
l2tp 1749
l2tp-access-profile 1751
layer2-unicast-replies 1752
lcp-renegotiation 1753
limit 1753
link (Address-Assignment Pools) 1754
local-gateway address 1754
liveness-detection 1755
lockout-timeout (L2TP Destination Lockout) 1756
local-server-group (DHCP Relay Agent Option) 1757
logical-bandwidth-policer 1758
logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces) 1758
logical-interface-policer 1759
logical-system (Diameter Peer) 1760
logical-system (Diameter Transport) 1761

Copyright © 2013, Juniper Networks, Inc.

logical-system (Tunnel Profile) 1761
logical-system-name (DHCP Local Server) 1762
logical-system-name (DHCP Relay Agent) 1763
logical-system-name (Static Subscribers) 1764
loss-priority (Dynamic Schedulers) 1765
loss-priority high then discard (Three-Color Policer) 1766
mac-address (DHCP Local Server) 1767
mac-address (DHCP Relay Agent) 1768
mac-address (Dynamic Access-Internal Routes) 1769
mac-validate (Dynamic IP Demux Interface) 1770
maintain-subscriber (Subscriber Management) 1771
managed-configuration (Dynamic Router Advertisement) 1772
mandatory 1772
map (Domain Map) 1773
match-direction 1774
match-direction (IP Reassembly Rule) 1774
mask (Domain Map) 1775
match-order (Dynamic Firewalls) 1775
max-advertisement-interval (Dynamic Router Advertisement) 1776
max-data-sessions-per-subscriber 1776
maximum-discovery-table-entries 1777
maximum-helper-restart-time 1777
maximum-lease-time 1778
maximum-sessions-per-tunnel 1778
max-outstanding-requests 1779
max-outstanding-requests (Gx-Plus) 1780
max-pending-accounting-stops (Access Profile) 1780
max-sessions (Dynamic PPPoE) 1781
max-sessions (PPPoE Service Name Tables) 1782
max-sessions (Tunnel Profile) 1783
max-sessions-vsa-ignore (Static and Dynamic Subscribers) 1784
max-withhold-time (Access Profile) 1785
medium (Tunnel Profile) 1785
method 1786
metric (Dynamic Access-Internal Routes) 1787
metric (Diameter Route) 1788
metric (Domain Map) 1788
min-advertisement-interval (Dynamic Router Advertisement) 1789
minimum-interval 1790
minimum-receive-interval 1791
mld (Dynamic Profiles) 1792
mobile-ip 1793
mode (Dynamic Profiles) 1794
multicast (Dynamic Routing Options) 1795
multicast-interception (Subscriber Secure Policy) 1796
multiplier 1797
nai 1798
name-server 1799
nas-identifier 1799

 Copyright © 2013, Juniper Networks, Inc.

nas-port-extended-format (Access Profile) 1800
nas-port-extended-format (Interfaces) 1802
nas-port-id-delimiter (Subscriber Management) 1803
nas-port-id-format (Subscriber Management) 1804
nas-port-method (Tunnel Profile) 1805
nas-port-options (RADIUS Options) 1806
nas-port-type (RADIUS Options) 1807
nas-port-type (Subscriber Management) 1809
neighbor (Associate with ANCP Access Identifier) 1810
neighbor (Define ANCP) 1811
neighbor-discovery-router-advertisement (Address-Assignment Pools) 1811
netbios-node-type 1812
network 1812
network-element (Diameter Base Protocol) 1813
next-hop (Dynamic Access-Internal Routes) 1814
next-hop-service 1815
no-accounting 1816
no-adaptation 1816
no-allow-snooped-clients 1817
no-bind-on-request (DHCP Relay Agent) 1818
no-keepalives (Dynamic Profiles) 1819
no-qos-adjust (Dynamic Routing Options) 1819
no-vlan-id-validate 1820
oam-on-svlan (Ethernet Interfaces) 1820
oif-map (Dynamic IGMP Interface) 1821
oif-map (Dynamic MLD Interface) 1821
on-demand-ip-address 1822
on-link (Dynamic Router Advertisement) 1822
option 1823
option-18 (Interface-ID for DHCPv6 Autosense VLANs) 1824
option-37 (Relay Agent Remote-ID for DHCPv6 Autosense VLANs) 1825
option-60 (DHCP Local Server) 1826
option-60 (DHCP Relay Agent) 1827
option-82 (Address-Assignment Pools) 1828
option-82 (DHCP Local Server Authentication) 1829
option-82 (DHCP Local Server Pool Matching) 1830
option-82 (DHCP Relay Agent) 1831
option-match 1832
option-number (DHCP Relay Agent Option) 1833
options (Access Profile) 1834
order 1835
order (Mobile IP) 1836
origin (Diameter Base Protocol) 1837
other-stateful-configuration (Dynamic Router Advertisement) 1837
output (Dynamic Service Sets) 1838
output-traffic-control-profile (Dynamic CoS Definition) 1839
output-vlan-map (Dynamic Interfaces) 1840
overhead-accounting (Dynamic Traffic Shaping) 1841
overrides (DHCP Local Server) 1842

Copyright © 2013, Juniper Networks, Inc.

overrides (DHCP Relay Agent) 1844
packet-triggered-subscribers 1845
packet-triggered-subscribers-partition 1846
padn (Domain Map) 1846
pap 1847
pap (Dynamic PPP) 1848
pap (L2TP) 1848
parse-direction (Domain Map) 1849
partition 1850
partition (Gx-Plus) 1850
partition (PTSP) 1851
passive (Dynamic IGMP Interface) 1852
passive (Dynamic MLD Interface) 1853
password (DHCP Local Server) 1854
password (DHCP Relay Agent) 1855
password (Static Subscribers) 1856
peak-burst-size 1857
peak-information-rate 1859
peer 1861
peer (Diameter Base Protocol) 1862
peer (Diameter Network Element) 1862
physical-interface-policer 1863
pic (M Series, MX Series, and T Series Routers) 1864
policer (Configuring) 1865
policy (Subscriber Secure Policy) 1867
policy-options 1868
pool (Address-Assignment Pools) 1869
pool (DHCP Local Server Overrides) 1870
pool (L2TP Service Interfaces) 1871
pool-match-order 1872
pop (Dynamic VLANs) 1873
port 1873
port (Diameter Peer) 1874
post-service-filter (Dynamic Service Sets) 1874
pp0 (Dynamic PPPoE) 1875
ppp (Group Profile) 1876
pppoe-options (Dynamic PPPoE) 1877
pppoe-tags (Adjustment Control Profiles) 1878
pppoe-underlying-options (Dynamic VLAN Interface Sets) 1879
pppoe-underlying-options (Static and Dynamic Subscribers) 1880
ppp-options 1881
ppp-options (Dynamic PPP) 1882
ppp-options (L2TP) 1883
ppp-subscriber-services 1884
precedence 1885
predefined-variable-defaults (Dynamic Profiles) 1886
preference (Subscriber Management) 1887
preference (Tunnel Profile) 1888
preferred-lifetime 1889

Copyright © 2013, Juniper Networks, Inc.

preferred-lifetime (Dynamic Router Advertisement) 1890
prefix (Address-Assignment Pools) 1890
prefix (DHCP Relay Agent) 1891
prefix (Dynamic Router Advertisement) 1893
pre-ietf-mode 1893
premium (Hierarchical Policer) 1894
priority (Diameter Peer) 1895
priority (Dynamic Schedulers) 1896
process-inform 1897
profile (Access) 1899
promiscuous-mode (Protocols IGMP) 1902
protocol (Dynamic Schedulers) 1903
protocol (Subscriber Secure Policy) 1903
protocols (Dynamic Profiles) 1904
provisioning-order 1906
proxy-arp 1906
proxy-mode 1907
ps0 (Pseudowire Subscriber Interfaces) 1908
pseudowire-service (Pseudowire Subscriber Interfaces) 1909
push (Dynamic VLANs) 1909
qos-adjust 1910
qos-adjust-adsl 1910
qos-adjust-adsl2 1911
qos-adjust-adsl2-plus 1911
qos-adjust-sds1 1912
qos-adjust-vds1 1912
qos-adjust-vds2 1913
qualified-bum-pruning-mode 1913
qualified-next-hop (Subscriber Management) 1914
radius (Access Profile) 1915
radius (Dynamic Profiles) 1917
radius-coa (Adjustment Control Profiles) 1918
radius-disconnect (DHCP Local Server) 1919
radius-flow-tap 1920
radius-options (Edit Access) 1921
radius-options (Interfaces) 1922
radius-server 1923
range (Address-Assignment Pools) 1924
rapid-commit (DHCPv6 Local Server) 1925
reachable-time (Dynamic Router Advertisement) 1926
rebalance-periodic (Aggregated Ethernet Subscriber Interfaces) 1926
realm (Diameter Origin) 1927
reconfigure (DHCP Local Server) 1928
registration-lifetime 1929
relay-agent-interface-id (DHCPv6 Relay Agent) 1930
relay-agent-interface-id (DHCPv6 Relay Agent Username) 1931
relay-agent-remote-id (DHCPv6 Relay Agent) 1932
relay-agent-interface-id (DHCP Local Server) 1933
relay-agent-remote-id (DHCP Local Server) 1934

Copyright © 2013, Juniper Networks, Inc.

relay-agent-subscriber-id (DHCP Local Server) 1935
relay-agent-subscriber-id (DHCPv6 Relay Agent) 1936
relay-option (DHCP Relay Agent) 1937
relay-option-82 1938
relay-server-group (DHCP Relay Agent Option) 1939
remote-gateway (Tunnel Profile) 1940
remote-id 1940
replace-ip-source-with 1941
replay-method 1942
report-interface-descriptions (Edit Access) 1943
request-rate 1943
retransmission-count-established (L2TP) 1944
retransmission-count-not-established (L2TP) 1945
retransmit-timer (Dynamic Router Advertisement) 1945
retry 1946
revert-interval 1947
revocation-required 1948
rewrite-rules (Dynamic CoS Interfaces) 1949
route (Access) 1950
route (Access Internal) 1951
route (Diameter Network Element) 1952
route-suppression (DHCP) 1953
router (Address-Assignment Pools) 1954
router-advertisement (Dynamic Profiles) 1954
routing-instance 1955
routing-instance (Diameter Peer) 1955
routing-instance (Diameter Transport) 1956
routing-instance (Tunnel Profile) 1956
routing-instance (PPPoE Service Name Tables) 1957
routing-instance-name (DHCP Local Server) 1958
routing-instance-name (DHCP Relay Agent) 1959
routing-instance-name (Static Subscribers) 1960
routing-instances 1961
routing-options (Dynamic Profiles) 1963
rpf-check (Dynamic Profiles) 1964
rule 1965
rule (IP Reassembly) 1966
rule-set 1967
rx-connect-speed-when-equal (L2TP LAC) 1967
scheduler (Dynamic Scheduler Maps) 1968
scheduler-map (Dynamic Traffic Shaping) 1968
scheduler-maps (Dynamic CoS Definition) 1969
schedulers (Dynamic CoS Definition) 1970
secret 1971
secret (Tunnel Profile) 1971
send-acct-status-on-config-change (Access Profile) 1972
send-release-on-delete (DHCP Relay Agent) 1973
server (Dynamic PPPoE) 1974
server-group 1975

 Copyright © 2013, Juniper Networks, Inc.

server-identifier (Address-Assignment Pools) 1976
service (Service Accounting) 1976
service (Dynamic Service Sets) 1977
service-device-pool (L2TP) 1978
service-device-pools (L2TP Service Interfaces) 1978
service-filter (Dynamic Service Sets) 1979
service-interface 1980
service-name-table 1981
service-profile (DHCP Local Server) 1982
service-profile (DHCP Relay Agent) 1983
service-set (Dynamic Service Sets) 1984
services (captive-portal-content-delivery) 1985
session-mode 1986
session-options 1987
shaping-rate (Dynamic Traffic Shaping and Scheduling) 1988
shared-name 1989
shared-secret 1989
short-cycle-protection (Static and Dynamic Subscribers) 1990
single-rate 1991
sip-server-address 1992
sip-server-domain-name 1992
source (Dynamic IGMP Interface) 1993
source (Dynamic MLD Interface) 1993
source-address 1994
source-address (Subscriber Secure Policy) 1994
source-count (Dynamic MLD Interface) 1995
source-gateway (Tunnel Profile) 1995
source-increment (Dynamic MLD Interface) 1996
source-ipv4-address 1996
source-port (Subscriber Secure Policy) 1997
spi 1998
ssm-map (Dynamic IGMP Interface) 1999
ssm-map (Dynamic MLD Interface) 1999
stacked-vlan-ranges (RADIUS Options) 2000
starts-with (DHCP Relay Agent Option) 2001
static (Dynamic IGMP Interface) 2002
static (Dynamic MLD Interface) 2003
static-subscribers 2004
statistics (Access Profile) 2005
strict (DHCP Local Server) 2006
strip-domain (Domain Map) 2007
subscriber-identification (PTSP) 2008
subscriber-management (Subscriber Management) 2009
subscriber-profile 2010
swap (Dynamic VLANs) 2010
tag (Access) 2011
tag (Dynamic Profiles) 2012
tag-protocol-id (Dynamic VLANs) 2012
target-logical-system (Domain Map) 2013

Copyright © 2013, Juniper Networks, Inc.

target-routing-instance (Domain Map) 2014
targeted-distribution (Static Interfaces over Aggregated Ethernet) 2015
term 2016
term (Captive Portal Content Delivery) 2017
terminate-code 2018
tftp-server 2019
then 2020
three-color-policer (Configuring) 2021
threshold (detection-time) 2022
threshold (transmit-interval) 2023
timeout (DHCP Local Server) 2024
timeout (RADIUS) 2025
timestamp-tolerance 2026
token (DHCP Local Server) 2027
tos-reflect (L2TP) 2028
trace (DHCP Local Server) 2029
trace (DHCP Relay Agent) 2030
traceoptions (General Authentication Service) 2031
traceoptions (ANCP) 2033
traceoptions (Captive Portal Content Delivery) 2035
traceoptions (DHCP) 2037
traceoptions (Diameter Base Protocol) 2039
traceoptions (L2TP) 2041
traceoptions (Mobile IP) 2045
traceoptions (PTSP) 2048
traceoptions (Static Subscribers) 2050
traceoptions (Subscriber Management) 2052
traceoptions (Subscriber Session Database Replication) 2054
traffic-control-profiles (Dynamic CoS Definition) 2055
transmit-interval 2056
transmit-rate (Dynamic Schedulers) 2057
transport (Diameter Base Protocol) 2058
transport (Diameter Peer) 2058
trigger (DHCP Local Server) 2059
trust-option-82 2060
tunnel (L2TP) 2061
tunnel (Tunnel Profile) 2062
tunnel-group 2063
tunnel-profile (Domain Map) 2064
tunnel-profile (L2TP Tunnel Switching) 2064
tunnel-profile (Tunnel Profile) 2065
tunnel-switch-profile (L2TP Tunnel Switching, Application) 2066
tunnel-switch-profile (L2TP Tunnel Switching, Definition) 2066
two-rate 2067
tx-address-change (L2TP LAC) 2068
tx-connect-speed-method (L2TP LAC) 2069
type (Tunnel Profile) 2070
t1-percentage 2071
t2-percentage 2072

Copyright © 2013, Juniper Networks, Inc.

uid 2073
underlying-interface (ANCP) 2073
underlying-interface (demux0) 2074
underlying-interface (Dynamic PPPoE) 2075
unit (Dynamic Demux Interface) 2076
unit (Dynamic Interface Sets) 2078
unit (Dynamic PPPoE) 2079
unit (Dynamic Traffic Shaping) 2081
unit (Dynamic Profiles Standard Interface) 2083
unnumbered-address (Dynamic Profiles) 2086
unnumbered-address (Dynamic PPPoE) 2087
untagged 2087
update-interval 2088
upstream-rate (Traffic Shaping) 2089
use-interface-description 2090
use-option-82 2091
use-primary (DHCP Local Server) 2092
use-primary (DHCP Relay Agent) 2093
user-group-profile 2094
user-prefix (DHCP Local Server) 2095
user-prefix (DHCP Relay Agent) 2096
user-prefix (Static Subscribers) 2097
username-include 2098
username-include (DHCP Local Server) 2099
username-include (DHCP Relay Agent) 2100
username-include (Static Subscribers) 2101
valid-lifetime 2102
valid-lifetime (Dynamic Router Advertisement) 2103
variables 2104
vendor-id 2105
vendor-specific-tags (Dynamic Traffic Shaping) 2105
version (BFD) 2106
version (Dynamic IGMP Interface) 2107
version (Dynamic MLD Interface) 2108
virtual-network 2109
vlan-id (Dynamic Profiles) 2110
vlan-id (Dynamic VLANs) 2111
vlan-model 2112
vlan-nas-port-stacked-format 2112
vlan-ranges (RADIUS Options) 2113
vlan-tag (Dynamic Classifiers) 2114
vlan-tag (Dynamic Rewrite Rules) 2115
vlan-tagging 2116
vlan-tags 2117
vrf-name (Duplicate Accounting) 2118
wait-for-acct-on-ack (Access Profile) 2118
weighted-load-balancing (L2TP LAC) 2119
wimax 2119
wins-server (Access) 2120

Copyright © 2013, Juniper Networks, Inc.

Part 20 Index

Index 2123
Index of Statements and Commands 2167

Copyright © 2013, Juniper Networks, Inc.

List of Figures

Part 1 Managing Access Networks

Chapter 1 **Subscriber Access Overview 3**

Figure 1: Subscriber Access Network Example 5
Figure 2: Subscriber Access Operation Flow 9
Figure 3: Subscriber Access Configuration Workflow 12

Part 2 AAA Service Framework for Subscriber Access

Chapter 2 **Configuring the AAA Service Framework for Subscriber Access 19**

Figure 4: Topology with Loss of Access to Accounting Server 33
Figure 5: DHCP Options Data Flow 82

Part 6 L2TP for Subscriber Access

Chapter 15 **L2TP for Subscriber Access Overview 379**

Figure 6: Typical L2TP Topology 380
Figure 7: Protocol Stacking for L2TP Subscribers in Pass-Through Mode 380
Figure 8: L2TP Tunnel Switching Network Topology 393
Figure 9: L2TP Tunnel Switching for Incoming Calls 394

Part 8 Mobile IP Access

Chapter 28 **Mobile IP Overview 557**

Figure 10: Mobile IP Network Without Reverse Tunneling 558
Figure 11: Mobile IP Network with Reverse Tunneling 559
Figure 12: Sample Mobile IP WiMAX Topology 567

Part 12 Subscriber Interfaces

Chapter 40 **Subscriber Interface Overview 757**

Figure 13: VLAN Subscriber Interfaces 758
Figure 14: IP Demux Subscriber Interface 759

Chapter 43 **Subscriber Interfaces over Aggregated Ethernet Overview 815**

Figure 15: Targeted Subscriber Links 819

Chapter 45 **Subscriber Interfaces over Aggregated Ethernet Examples 835**

Figure 16: Multicast Traffic Separation Using OIF Mapping 846

Chapter 49 **MPLS Pseudowire Subscriber Interfaces Overview 923**

Figure 17: MPLS Access Network with Subscriber Management Support 924
Figure 18: Pseudowire Subscriber Interface Protocol Stack 924

Copyright © 2013, Juniper Networks, Inc.

Figure 19: MPLS Pseudowire Subscriber Interface Two-Level Scheduler Configuration 926
Figure 20: Three-Level Scheduling Hierarchy Case 1: Pseudowire Service Logical Interfaces over a Transport Logical Interface 928
Figure 21: Three-Level Scheduling Hierarchy Case 2: Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set 929
Figure 22: Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces-Deployment Scenario 929

Part 13 Class of Service for Subscriber Access

Chapter 51 **Dynamic CoS for Subscriber Access Overview 949**

Figure 23: Two-Level Hierarchical Scheduling 957
Figure 24: Three-Level Hierarchical Scheduling: Logical Interfaces at Level 3 with Underlying Logical Interfaces at Level 2 959
Figure 25: Three-Level Hierarchical Scheduling: Logical Interfaces at Level 2 with Interface Sets at Level 3 959
Figure 26: Logical Interfaces at Level 2 and Interface Sets at Level 3 960
Figure 27: Logical Interfaces at Level 3 and Underlying Logical Interfaces at Level 2 961

Chapter 56 **Interface Solutions for Dynamic CoS Overview 1013**

Figure 28: CoS Configuration for L2TP LAC Topology 1016
Figure 29: Processing of CoS Parameters in an L2TP LNS Inline Service 1018

Chapter 59 **Bandwidth Management for Dynamic CoS Overview 1069**

Figure 30: Sample Burst Shaping Rates 1070
Figure 31: Distribution of Queues on the 30-Gigabit Ethernet Queuing MPC with One MIC 1077
Figure 32: Distribution of Queues on the 30-Gigabit Ethernet Queuing MPC with Two MICs 1078
Figure 33: Distribution of Queues on the 60-Gigabit Ethernet Enhanced Queuing MPC 1078
Figure 34: Scheduler Node and Queues with Adjusted Shaping Rates 1081
Figure 35: Queue with Adjusted Shaping Rate 1081

Chapter 61 **Bandwidth Management for Dynamic CoS Examples 1111**

Figure 36: Sample Network Topology for Downstream Traffic 1111

Part 15 Dynamic Firewall Filters, Service Sets and HTTP Redirect for Subscriber Access

Chapter 68 **Firewall Filter Examples 1199**

Figure 37: Logical Flow Example for Filter Bypass Processing 1211

Part 16 Subscriber Secure Policy Traffic Mirroring

Chapter 73 **Overview of Subscriber Secure Policy Using RADIUS 1247**

Figure 38: RADIUS-Initiated Subscriber Secure Policy Architecture 1248
Figure 39: RADIUS-Initiated Traffic Mirroring Interfaces 1250
Figure 40: RADIUS-Initiated Subscriber Secure Policy Model at Login 1251

 Copyright © 2013, Juniper Networks, Inc.

Figure 41: RADIUS-Initiated Subscriber Secure Policy Model After Login 1252
Figure 42: Mirrored Packet Header and Payload 1256
Figure 43: 4-Byte Format of VSA 26-59 1258
Figure 44: 8-Byte Format of VSA 26-59 1259

Chapter 75 **Overview of Subscriber Secure Policy Using DTCP 1267**

Figure 45: DTCP-Initiated Subscriber Secure Policy Architecture 1268
Figure 46: DTCP-Initiated Traffic Mirroring Interfaces 1269
Figure 47: DTCP-Initiated Subscriber Secure Policy Model 1271
Figure 48: Mirrored Packet Header 1278

Part 17 Protocols for Subscriber Access

Chapter 79 **ANCP Overview 1311**

Figure 49: Sample ANCP Topology Without Interface Sets (1:1 and N:1 Model) 1322
Figure 50: Sample ANCP Topology with Interface Sets (N:1 Model) 1324
Figure 51: Sample ANCP Topology with Interface Sets (1:1 Model) 1326

Chapter 80 **Configuring ANCP 1331**

Figure 52: N:1 ANCP Topology with Interface Sets and VLAN Demux Interface over Aggregated Ethernet 1350

Copyright © 2013, Juniper Networks, Inc.

List of Tables

About This Guide lxvii

Table 1: Notice Icons lxxi
Table 2: Text and Syntax Conventions lxxi

Part 1 Managing Access Networks

Chapter 1 **Subscriber Access Overview 3**

Table 3: Subscriber Access Terms and Acronyms 4

Part 2 AAA Service Framework for Subscriber Access

Chapter 2 **Configuring the AAA Service Framework for Subscriber Access 19**

Table 4: RADIUS Attributes and VSAs Used for Per-Subscriber Session Accounting 27
Table 5: Duplicate RADIUS Accounting Reporting 28
Table 6: Juniper Networks VSAs Used for Per-Service Session Accounting 36
Table 7: RADIUS NAS-Port-Type Values 46
Table 8: Attributes That Can Be Ignored in RADIUS Access-Accept Messages 58
Table 9: Attributes That Can Be Excluded from RADIUS Messages 59
Table 10: Unsupported Opaque DHCP Options 84
Table 11: DHCP Options Description 85
Table 12: Identification Attributes 87
Table 13: Session Attributes 88
Table 14: Supported RADIUS IETF Attributes 91
Table 15: Supported Juniper Networks VSAs 99
Table 16: DSL Forum VSAs 108
Table 17: AAA Access Messages—Supported RADIUS Attributes and Juniper Networks VSAs 110
Table 18: AAA Accounting Messages—Supported RADIUS Attributes and Juniper Networks VSAs 115
Table 19: DSL Forum VSAs—Supported RADIUS Messages 119
Table 20: Error-Cause Codes (RADIUS Attribute 101) 120
Table 21: Supported RADIUS Acct-Terminate-Cause Codes 121
Table 22: Default AAA Mappings 123
Table 23: Default DHCP Mappings 124
Table 24: Default L2TP Mappings 124
Table 25: Default PPP Mappings 141

Chapter 3 **Configuring Address-Assignment Pools for Subscriber Access 167**

Table 26: DHCP Attributes 173

Table 27: DHCPv6 Attributes 174

Chapter 4 **Configuring Domain Maps for Subscriber Access 179**

Table 28: Domain Map Options and Parameters 180
Table 29: Precedence Rules for Applying Access Profiles 182
Table 30: Precedence Rules for Applying Dynamic Profiles 183
Table 31: Precedence Rules for Determining the Address Pool to Use 184

Part 3 DHCP Local Server for Subscriber Access

Chapter 6 **DHCP Local Server Overview 199**

Table 32: Information in Authentication Grant 202
Table 33: RADIUS Attributes and VSAs for DHCPv6 Local Server 204

Chapter 7 **Configuring DHCP Local Server 211**

Table 34: Actions for DHCP Local Server Snooped Packets 224
Table 35: Action Taken for Events That Occur During a Reconfiguration 244

Part 4 DHCP Relay Agent for Subscriber Access

Chapter 10 **Configuring DHCP Relay Agent 285**

Table 36: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled 307
Table 37: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Disabled 307
Table 38: Actions for Snooped BOOTREPLY Packets 307
Table 39: DHCP Relay Agent Option 82 Value for Auto Logout 310

Part 6 L2TP for Subscriber Access

Chapter 15 **L2TP for Subscriber Access Overview 379**

Table 40: L2TP Terms 382
Table 41: Cause of CDN Message 395
Table 42: Cause of StopCCN Message 395
Table 43: LAC, LNS, and LTS Actions Taken for Switched Tunnels in Response to Administrative clear Commands 396
Table 44: Default Action for Handling L2TP AVPs at the Switching Boundary . . 397

Chapter 16 **Configuring L2TP for Subscriber Access 401**

Table 45: SNMP Counters for L2TP Statistics 416
Table 46: VSA and Standard RADIUS Attribute Names, Order, and Values Required for Example 441

Part 7 Diameter Base Protocol and Applications for Subscriber Access

Chapter 17 **Diameter Base Protocol Overview 455**

Table 47: Diameter Messages and Diameter Applications 457
Table 48: Standard Diameter AVPs 462
Table 49: Juniper Networks Diameter AVPs 466
Table 50: Tekelec Diameter AVPs 469

Chapter 25 **Configuring Packet-Triggered Subscriber Services 529**

 Copyright © 2013, Juniper Networks, Inc.

Table 51: PTSP Match Conditions 532
Table 52: PTSP Actions 533
Table 53: PTSP Forward Rule Match Conditions 535

Chapter 26 **Gx-Plus Overview 541**

Table 54: Differences Between Gx-Plus and Junos OS Terminology 541
Table 55: Router Events, Router Actions, and PCRF Actions 546

Part 8 Mobile IP Access

Chapter 28 **Mobile IP Overview 557**

Table 56: Juniper Networks VSAs Used by Mobile IP 561
Table 57: WiMAX Forum VSAs used by Mobile IP 566

Part 9 ATM for Subscriber Access

Chapter 30 **ATM for Subscriber Access Overview 581**

Table 58: Encapsulation Types for Supported ATM Subscriber Access Configurations 585

Part 10 Dynamic Profiles for Access and Services

Chapter 33 **Dynamic Profiles Overview 637**

Table 59: Junos OS Predefined Variables and Definitions 642
Table 60: RADIUS Attributes and Corresponding Junos OS Predefined Variables 660
Table 61: Operators and Functions 668
Table 62: Expression Errors and Actions 669
Table 63: Types of Variables Supported in Dynamic Profiles 671
Table 64: Default Values and Expressions Supported in Dynamic Profiles 671

Chapter 34 **Configuring Dynamic Profiles 673**

Table 65: Expression Examples 678

Part 11 Dynamic VLANs

Chapter 37 **Dynamic VLAN Overview 701**

Table 66: CoS Adjustment in Dynamic Profiles for ACI Interface Sets and ACI-Based Subscriber Interfaces 708
Table 67: Ethernet OAM Support for Service VLANs Terms and Acronyms 710

Part 12 Subscriber Interfaces

Chapter 40 **Subscriber Interface Overview 757**

Table 68: Comparison of MAC Address Validation Modes 763

Chapter 43 **Subscriber Interfaces over Aggregated Ethernet Overview 815**

Table 69: Features Supported with Static or Dynamic Demux Subscriber Interfaces 818

Chapter 46 **Dynamic PPPoE Subscriber Interfaces Overview 883**

Table 70: Sample PPPoE Maximum Session Values During Subscriber Login . . 892

Copyright © 2013, Juniper Networks, Inc.

Chapter 48 Dynamic PPPoE Subscriber Interfaces Examples 917
Table 71: Dynamic PPPoE Subscriber Interface Creation Based on PPPoE Client Request Values 921
Chapter 49 MPLS Pseudowire Subscriber Interfaces Overview 923
Table 72: Two-Level Hierarchical Scheduling-Interface Hierarchy Versus Scheduling Nodes 926
Table 73: Three-Level Hierarchical Scheduling-Interface Hierarchy Versus CoS Scheduling Node Levels 927

Part 13 Class of Service for Subscriber Access

Chapter 51 Dynamic CoS for Subscriber Access Overview 949
Table 74: Hardware Required for Dynamic CoS Configurations 951
Table 75: IP Demux Classification Rules 954
Table 76: IP Demux Rewrite Rules 954
Table 77: L2TP Classification Rules 955
Table 78: L2TP LAC Rewrite Rules 955
Table 79: Two-Level Hierarchical Scheduling-Interface Hierarchy Versus Scheduling Nodes 957
Table 80: Three-Level Hierarchical Scheduling-Interface Hierarchy Versus CoS Scheduling Node Levels 958
Chapter 54 RADIUS and Dynamic CoS Overview 989
Table 81: CoS Predefined Variables for Scheduler Map and Traffic Shaping . . . 990
Table 82: CoS Predefined Variables for Scheduling and Queuing 991
Table 83: CoS Services and Variables 994
Chapter 55 Configuring RADIUS for Dynamic CoS 999
Table 84: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets 1006
Chapter 56 Interface Solutions for Dynamic CoS Overview 1013
Table 85: Ingress LAC Tunnel Classifier Options 1016
Table 86: Sample Result 1017
Table 87: Hardware Requirements for L2TP LNS Inline Services 1019
Chapter 58 Dynamic CoS for Subscriber Access Examples 1031
Table 88: Initial Scheduler Map and Shaping Values at Subscriber Login 1034
Table 89: Initial CoS Values for the Voice Scheduler at Subscriber Login 1035
Table 90: Initial CoS Values for the Data Scheduler at Subscriber Login 1035
Table 91: Upgraded CoS Values for the Video Service 1038
Table 92: Upgraded CoS Values for the Video Scheduler 1038
Table 93: Initial CoS Values for the Expedited Forwarding Scheduler at Subscriber Login 1039
Table 94: Initial CoS Values for the Best Effort Scheduler at Subscriber Login . . 1039
Table 95: Scheduler Per Logical Interface Mapping 1045
Table 96: Scheduler per Underlying Interface Mapping 1048
Table 97: Scheduler per Logical Interface with Interface Set Mapping 1050
Chapter 59 Bandwidth Management for Dynamic CoS Overview 1069

 Copyright © 2013, Juniper Networks, Inc.

Table 98: Dedicated Queues for MIC and MPC Interfaces 1076
Table 99: Terms and Definitions for Shaping Characteristics 1084
Table 100: Adjustment Control Profile Applications and Algorithms 1085

Chapter 61 **Bandwidth Management for Dynamic CoS Examples 1111**

Table 101: Initial Shaping Values at Subscriber Login 1112
Table 102: Initial Shaping Values at Subscriber Login 1114

Part 15 Dynamic Firewall Filters, Service Sets and HTTP Redirect for Subscriber Access

Chapter 63 **Dynamic Firewall Filters and Service Sets Overview 1131**

Table 103: Ascend-Data-Filter Attribute Fields 1138
Table 104: Enhanced Network Services Mode and Firewall Filter Use Case Determination 1142

Chapter 64 **Parameterized Filters 1153**

Table 105: Dynamic Profiles and UID Substitution Comparison 1158

Chapter 66 **Configuring Fast Update Filters 1183**

Table 106: Fast Update Filter Match Conditions 1186
Table 107: Fast Update Filter Actions and Action Modifiers 1187

Chapter 68 **Firewall Filter Examples 1199**

Table 108: Ascend-Data-Filter Rule 1207

Part 16 Subscriber Secure Policy Traffic Mirroring

Chapter 73 **Overview of Subscriber Secure Policy Using RADIUS 1247**

Table 109: RADIUS-Initiated Subscriber Secure Policy Functions and Components 1248
Table 110: RADIUS-Initiated Traffic Mirroring Interfaces 1250
Table 111: RADIUS-Based Mirroring Attributes 1254
Table 112: RADIUS Attributes Used in CoA Messages to Identify Subscribers for Traffic Mirroring 1255
Table 113: Mirrored Packet Header and Payload Field Descriptions 1256

Chapter 75 **Overview of Subscriber Secure Policy Using DTCP 1267**

Table 114: DTCP-Initiated Subscriber Secure Policy Functions and Components 1268
Table 115: DTCP-Initiated Traffic Mirroring Interfaces 1270
Table 116: DTCP Mirroring Triggers for Use in ADD Messages 1274
Table 117: Mirrored Packet Header and Payload Field Descriptions 1278

Chapter 78 **Reporting Intercept-Related Events for Subscriber Secure Policy 1305**

Table 118: Subscriber Secure Policy SNMPv3 Traps for LAES Messages 1306

Part 17 Protocols for Subscriber Access

Chapter 79 **ANCP Overview 1311**

Table 119: ACI Mapping by Interface Type for the ANCP 1:1 Model 1320

Table 120: ACI Mapping by Interface Set Type for the ANCP N:1 Access Model 1320
Table 121: Mapping ANCP DSL Attributes to Juniper Networks DSL VSAs and DSL Forum VSAs 1328

Chapter 80 **Configuring ANCP 1331**

Table 122: Configuration Components used in ANCP N:1 Topology Example with Interface Sets 1351

Part 19 Complete Configuration Statement Hierarchy and Summary of Statements for Subscriber Access

Chapter 86 **Subscriber Access Configuration Statements 1413**

Table 123: Bandwidth Limits and Token Rates 1486

Copyright © 2013, Juniper Networks, Inc.

About This Guide

This preface provides the following guidelines for using the *Junos® OS Junos OS Subscriber Management Feature Guide*:

- JUNOS Documentation and Release Notes on page lxvii
- Objectives on page lxviii
- Audience on page lxviii
- Supported Routing Platforms on page lxix
- Using the Indexes on page lxix
- Using the Examples in This Manual on page lxix
- Documentation Conventions on page lxx
- Documentation Feedback on page lxxii
- Requesting Technical Support on page lxxii

JUNOS Documentation and Release Notes

For a list of related JUNOS documentation, see http://www.juniper.net/techpubs/software/junos/.

If the information in the latest release notes differs from the information in the documentation, follow the *JUNOS 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/.

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using the Junos operating system (Junos OS) and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using Junos OS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at http://www.juniper.net/books.

Objectives

This guide provides an overview of the subscriber access management features of the Junos OS and describes how to configure and manage remote subscriber access on the routing platform.

NOTE: **For additional information about Junos OS—either corrections to or information that might have been omitted from this guide—see the software release notes at http://www.juniper.net.**

Audience

This guide is designed for network administrators who are configuring and monitoring Juniper Networks MX Series Ethernet Services Routers.

To use this guide, you need a broad understanding of networks in general, the Internet in particular, networking principles, and network configuration. You must also be familiar with one or more of the following Internet routing protocols:

- Border Gateway Protocol (BGP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Internet Control Message Protocol (ICMP) router discovery
- Internet Group Management Protocol (IGMP)
- Multiprotocol Label Switching (MPLS)
- Open Shortest Path First (OSPF)
- Protocol-Independent Multicast (PIM)
- Resource Reservation Protocol (RSVP)
- Routing Information Protocol (RIP)
- Simple Network Management Protocol (SNMP)

Personnel operating the equipment must be trained and competent; must not conduct themselves in a careless, willfully negligent, or hostile manner; and must abide by the instructions provided by the documentation.

Copyright © 2013, Juniper Networks, Inc.

Supported Routing Platforms

For the features described in this manual, the Junos OS currently supports the following routing platforms:

- MX Series routers

Using the Indexes

This reference contains two indexes: a complete index that includes topic entries, and an index of statements and commands only.

In the index of statements and commands, an entry refers to a statement summary section only. In the complete index, the entry for a configuration statement or command contains at least two parts:

- The primary entry refers to the statement summary section.
- The secondary entry, usage guidelines, refers to the section in a configuration guidelines chapter that describes how to use the statement or command.

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

```
system {
  scripts {
    commit {
      file ex-script.xsl;
    }
  }
```

```
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Copyright © 2013, Juniper Networks, Inc.

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.

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the**configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration 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)	Enclose optional keywords or variables.	**stub <default-metric** *metric***>;**

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

Convention	Description	Examples
| (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)	Enclose a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identify 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		
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 send your comments to **techpubs-comments@juniper.net**, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/ . If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release 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 JNASC support contract,

Copyright © 2013, Juniper Networks, Inc.

or are covered under warranty, and need postsales 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:

- Find CSC offerings: http://www.juniper.net/customers/support/
- 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: https://www.juniper.net/alerts/
- 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, visit us at http://www.juniper.net/support/requesting-support.html

Copyright © 2013, Juniper Networks, Inc.

PART 1

Managing Access Networks

- Subscriber Access Overview on page 3

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 1

Subscriber Access Overview

- Subscriber Access Overview on page 3
- Subscriber Access Environment on page 4
- Relationship Between Subscribers and Interfaces in an Access Network on page 5
- Subscriber Access Support Considerations on page 5
- Subscriber Access Licensing Overview on page 6
- Subscriber Management Unified ISSU Support on page 7
- Verifying and Monitoring Subscriber Management Unified ISSU State on page 8
- Subscriber Access Operation Flow Using DHCP Relay on page 8
- Subscriber Activation and Service Management in an Access Network on page 9
- Configuring Subscriber Access on page 10
- Collecting Subscriber Access Logs Before Contacting Juniper Technical Support on page 13
- Compressing Troubleshooting Logs from /var/logs to Send to Juniper Technical Support on page 15

Subscriber Access Overview

The Juniper Networks Junos OS subscriber access feature provides subscriber access, authentication, and service creation, activation, and deactivation. You can also collect accounting information and statistics for subscriber service sessions.

The subscriber access feature supports both CLI and AAA-based configuration (such as RADIUS) for subscribers. Access and services start when the router receives a message from a client (such as a DHCP discover message). For RADIUS clients, RADIUS Access-Accept messages and Change-of-Authorization-Request (CoA-Request) messages can create, modify, and delete subscriber sessions as well as activate and deactivate service sessions. You can use CLI commands to create a dynamic profile, which acts as a template of user attributes.

A subscriber service is based on the combination of a defined dynamic profile and attributes configured through authentication. Dynamic profiles can include dynamic firewall filters, class-of-service (CoS) settings, and protocol (IGMP) settings that define access limits for subscribers and the scope of a service granted to the subscriber after access is obtained.

Copyright © 2013, Juniper Networks, Inc.

The subscriber access feature provides the following convenience and flexibility to service providers and subscribers:

- Service providers can separate services and access technology and eliminate unprofitable flat-rate billing. They gain the ability to efficiently design, manage, and deliver services that subscribers want, and then bill subscribers based on connect time, bandwidth, and the actual service used.
- Subscribers benefit by gaining access to multiple simultaneous services. Depending on the service provider configuration, subscribers can dynamically connect to and disconnect from various services when they want and for however long they want. Subscribers can be billed based on the service level and usage, rather than being charged a set rate regardless of usage.

Subscriber Access Terms and Acronyms

Table 3 on page 4 defines terms and acronyms that are used in this discussion of subscriber access.

Table 3: Subscriber Access Terms and Acronyms

Term	Definition
AAA method for subscriber authentication	The AAA method that uses authentication (for example, including RADIUS VSAs in the Access-Accept packet) to verify a subscriber and activate a service when the subscriber logs in.
Dynamic profile	A template that defines a set of characteristics that are combined with authorization attributes and are dynamically assigned to static interfaces to provide dynamic subscriber access and services for broadband applications.
RADIUS CoA method	The method that uses RADIUS CoA-Request messages and VSAs to activate a service for a subscriber that is already logged in.
Subscriber access technology	The technology used by a subscriber to access services (for example, DHCP).

Related Documentation

- Subscriber Access Environment on page 4
- Subscriber Access Licensing Overview on page 6
- Subscriber Access Operation Flow Using DHCP Relay on page 8
- Configuring Subscriber Access on page 10

Subscriber Access Environment

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

Copyright © 2013, Juniper Networks, Inc.

The subscriber access technologies include:

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

The subscriber authentication protocols include the RADIUS server.

Figure 1 on page 5 shows an example of a basic subscriber access network.

Figure 1: Subscriber Access Network Example

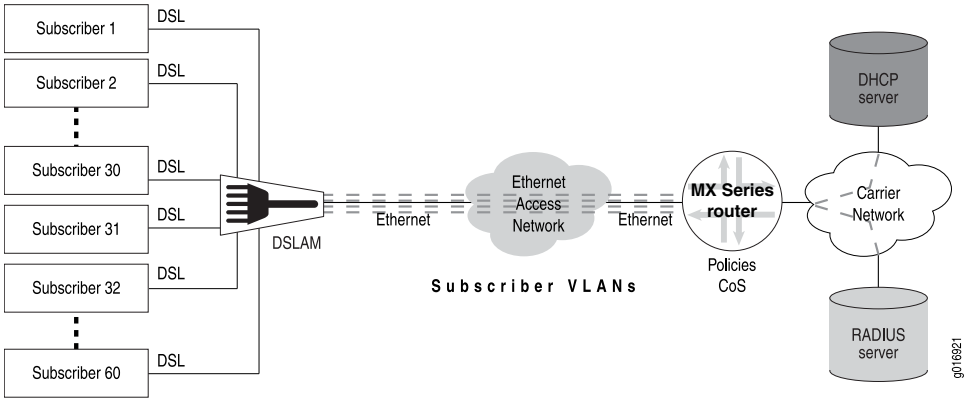

Related Documentation

- Subscriber Access Overview on page 3

Relationship Between Subscribers and Interfaces in an Access Network

To the router, a subscriber is an authenticated user. This release supports configurations of only one subscriber per logical interface. However, a subscriber can be either one authenticated client or a group of clients on a single, authenticated VLAN.

Related Documentation

- Subscriber Interface Overview on page 757

Subscriber Access Support Considerations

The subscriber access feature is limited to MX Series 3D Universal Edge Routers and the interfaces you can use when configuring dynamic profiles.

Platform Support

Even though many statements appear in the CLI for various other platforms, Juniper Networks supports subscriber access DHCP configuration on MX Series routers only. In addition, PPPoE configuration is currently supported on MX Series routers, M120 routers, and M320 routers.

Copyright © 2013, Juniper Networks, Inc.

Interface Support

You can use dynamic profiles to configure statically created interfaces and also to create and configure interfaces dynamically. Subscriber interfaces support IPv4 and IPv6 addressing.

To identify subscribers statically, you can reference a static VLAN interface in a dynamic profile. To identify subscribers dynamically, you create variables for demux interfaces that are dynamically created when subscribers log in.

The subscriber access feature supports the following device types:

- GE -- Gigabit Ethernet
- XE -- 10–Gigabit Ethernet
- AE -- Aggregated Ethernet

DPC Support

Certain subscriber management features require the use of specific dense port concentrators (DPCs) on the MX series router. For a list of the MX series DPCs and the features they support, see the *MX Series Interface Module Reference*.

Routing Engine Support

Subscriber management features require the use of either the RE-S-2000 or RE-S-1800 routing engine. The RE-S-1300 routing engine is not supported for use with subscriber management configurations.

Logical System Support

Subscriber management is supported in the default logical system only.

Related Documentation

- Relationship Between Subscribers and Interfaces in an Access Network on page 5
- Configuring Subscriber Access on page 10

Subscriber Access Licensing Overview

To enable some Juniper Networks Junos OS features or router scaling levels, you might have to purchase, install, and manage separate software license packs. The presence on the router of the appropriate software license keys (passwords) determines whether you can configure and use certain features or configure a feature to a predetermined scale.

Related Documentation

- Configuring the Router to Strictly Enforce the Subscriber Scaling License on page 235
- For information about installing and managing Junos OS licenses, see the *Installation and Upgrade Guide*

Copyright © 2013, Juniper Networks, Inc.

Subscriber Management Unified ISSU Support

The unified in-service software upgrade (unified ISSU) feature supports both the DHCP access model and the PPPoE access model used by subscriber management. This support ensures that the router preserves active DHCP and PPPoE subscriber sessions and session services after a unified ISSU has completed.

The *Junos OS High Availability Library for Routing Devices* describes the supported platforms and modules, CLI statements, and procedures you use to configure and initiate unified ISSU. You can use the **issu** flag with the **traceoptions** statement to trace subscriber management unified ISSU events. You can also use the **show system subscriber-management summary** command to display information about the unified ISSU state.

This overview describes specific considerations for unified ISSU support of the DHCP and PPPoE access models for subscriber management, and covers the following topics:

- Unified ISSU Support for DHCP Access Model on page 7
- Unified ISSU Support for PPPoE Access Model on page 7

Unified ISSU Support for DHCP Access Model

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

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

Unified ISSU Support for PPPoE Access Model

Unified ISSU supports the subscriber management PPPoE access model for static and dynamic PPPoE access, and includes the following features:

- Terminated, non-tunneled PPPoE connections configured with static or dynamic PPP logical interfaces and static or dynamic underlying interfaces
- Subscriber services on single-link PPP interfaces
- Preservation of statistics for accounting, filter, and CoS on MPC/MIC interfaces

  Accounting statistics for PPPoE subscribers are *not* preserved after a unified ISSU on Enhanced Intelligent Queuing 2 (IQ2E) PICs on M120 and M320 routers.

Unified ISSU for the subscriber management PPPoE access model *does not support* Multilink Point-to-Point Protocol (MLPPP) bundle interfaces. MLPPP bundle interfaces require the use of an Adaptive Services PIC or Multiservices PIC to provide PPP subscriber services. These PICs do not support unified ISSU.

Related Documentation

- Verifying and Monitoring Subscriber Management Unified ISSU State on page 8

- *Unified ISSU Support on MX Series 3D Universal Edge Routers* in *Unified ISSU System Requirements*

Copyright © 2013, Juniper Networks, Inc.

- For information about unified ISSU, see the *Junos OS High Availability Library for Routing Devices*

Verifying and Monitoring Subscriber Management Unified ISSU State

Purpose Display the state of unified ISSU for subscriber management features.

Action The first example indicates that control plane quiescing as part of unified ISSU is not in progress (for example, unified ISSU has not been started, has already completed, or control plane queiscing has not started). The second example shows that unified ISSU is in progress and that a participating subscriber management daemon requires 198 seconds to quiesce the control plane.

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  IDLE
    ISSU State           IDLE
    ISSU Wait            0
```

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  DAEMON_ISSU_PREPARE
    ISSU State           PREPARE
    ISSU Wait            198
```

Related Documentation

- Subscriber Management Unified ISSU Support on page 7
- See the “Unified ISSU” chapter in the *Junos OS High Availability Library for Routing Devices*

Subscriber Access Operation Flow Using DHCP Relay

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

Figure 2 on page 9 shows the flow of operations that occurs when the router is using DHCP relay to enable access for a subscriber.

 Copyright © 2013, Juniper Networks, Inc.

Figure 2: Subscriber Access Operation Flow

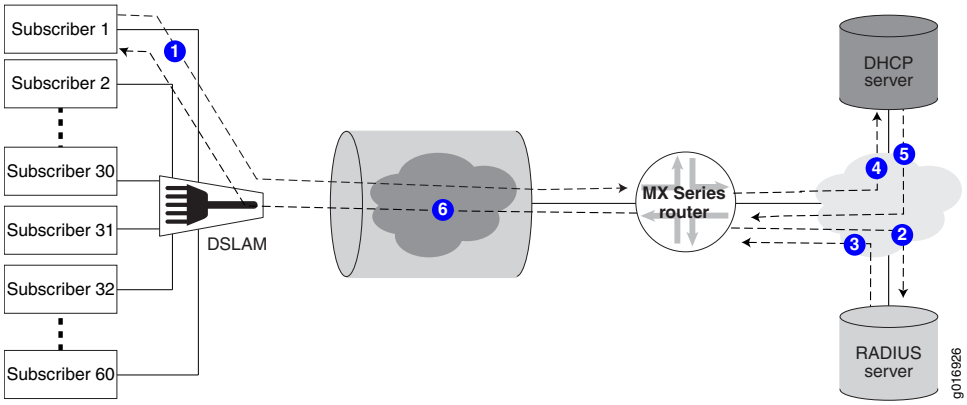

The following general sequence occurs during access configuration for a DHCP client:

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

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

Related Documentation

- Subscriber Access Overview on page 3
- Configuring Subscriber Access on page 10

Subscriber Activation and Service Management in an Access Network

The subscriber access feature uses dynamic profiles to activate subscribers and manage services.

A dynamic profile is a set of characteristics, defined in a template, that the router uses to provide dynamic subscriber access and services.

By using dynamic profiles you can:

- Define access for your network
- Define different service levels for subscribers
- Preprovision services that you can activate later

Using AAA-based login (RADIUS-based login or RADIUS CoA) you can:

- Provide subscribers with dynamic activation and deactivation based on service selection

Copyright © 2013, Juniper Networks, Inc.

- Provide greater flexibility and efficient management for a large number of subscribers and services

Components of a Dynamic Profile

You can use dynamic profiles to define various router components for subscriber access.

These components include the following:

- Dynamic firewall filters—Includes input and output filters to enforce rules that define whether to permit or deny packets that are transmitting an interface on the router. To apply dynamic firewall filters to the subscriber interface, you configure static input and output firewall filters and reference those filters in dynamic profiles.
- Dynamic Class of Service (CoS)—Includes CoS values that define a service for a subscriber. For example, you can configure the shaping rate for traffic in a video service by referencing CoS statements in a dynamic profile.
- Dynamic signaling protocol—Includes dynamic IGMP configuration for host to router signaling for IPv4 to support IP multicasting.

Router Predefined Variables Used by Dynamic Profiles

The router contains many predefined variables. These variables enable dynamic association of certain interface-specific values to incoming subscriber requests. You must specify these predefined variables in certain statements within a dynamic profile. When a client accesses the router, the dynamic profile configuration replaces the predefined variable with the actual data from an incoming client data packet and configuration (local and RADIUS).

Related Documentation

- Dynamic Profiles Overview on page 638
- Subscriber Interface Overview on page 757
- Junos OS Predefined Variables on page 642

Configuring Subscriber Access

To configure subscriber access:

1. Configure the client access protocol.
   - Configure DHCP local server.

     See "Extended DHCP Local Server Overview" on page 200.
   - Configure DHCP relay.

     See "Extended DHCP Relay Agent Overview" on page 276.
   - Configure PPP.

     See the "Configuring Logical Interface Properties" and "Configuring Point-to-Point Protocol over Ethernet" chapters of the *Junos OS Network Interfaces Library for Routing Devices*.

 Copyright © 2013, Juniper Networks, Inc.

2. Configure subscriber authentication, accounting, and addressing.
   a. Configure RADIUS:
      1. Specify the RADIUS servers.

         See "Specifying RADIUS Authentication and Accounting Servers for Subscriber Access" on page 37.

      2. Specify any optional server attributes.

         See "Configuring RADIUS Server Options for Subscriber Access" on page 47.

      3. (Optional) Configure the CoA feature for the RADIUS dynamic-request server to change or deactivate the service after login.

         See "Configuring RADIUS-Initiated Dynamic Request Support" on page 90.

      4. Configure subscriber accounting (RADIUS accounting).

         See "Configuring Per-Subscriber Session Accounting" on page 29.

   b. Configure addressing:
      - See "Configuring Address-Assignment Pools" on page 168.

3. Create and manage dynamic profiles for access and service.
   a. Configure a basic dynamic profile.

      See "Configuring a Basic Dynamic Profile" on page 673.

      See "Example: Minimum PPPoE Dynamic Profile" on page 697

   b. Configure a dynamic profile for access.

      See "Configuring a Dynamic Profile for Client Access" on page 679.

   c. Configure a dynamic profile for services.

      See "Configuring a Dynamic Profile for Various Levels of Services" on page 681.

   d. Configure a default subscriber service.

      See "Configuring a Default Subscriber Service" on page 1127.

   e. Configure the static subscriber interfaces to be referenced in the dynamic profile.

      See "Configuring a Subscriber Interface with a Static VLAN Interface" on page 766.

   f. Specify the interface-name and unit variables that the router uses to dynamically associate to a subscriber's incoming interface.

      See "Associating Dynamic Profiles with Statically Created Interfaces" on page 768.

   g. Add, modify, or delete dynamic profile values to manage subscriber access and services.

      See "Modifying Dynamic Profiles with Versioning Disabled" on page 684.

The router dynamically activates or modifies the subscriber service using the RADIUS configuration.

Copyright © 2013, Juniper Networks, Inc.

- When the subscriber logs in, the router dynamically activates the service.

  See "Dynamic Service Activation During Login Overview" on page 86.

- If RADIUS CoA has been configured, the router can dynamically modify the service for a subscriber.

  See "RADIUS-Initiated Change of Authorization (CoA) Overview" on page 87.

Figure 3 on page 12 shows the configuration sequence you perform for DHCP-based subscriber access. It also shows the dynamic configuration performed by the router.

Figure 3: Subscriber Access Configuration Workflow

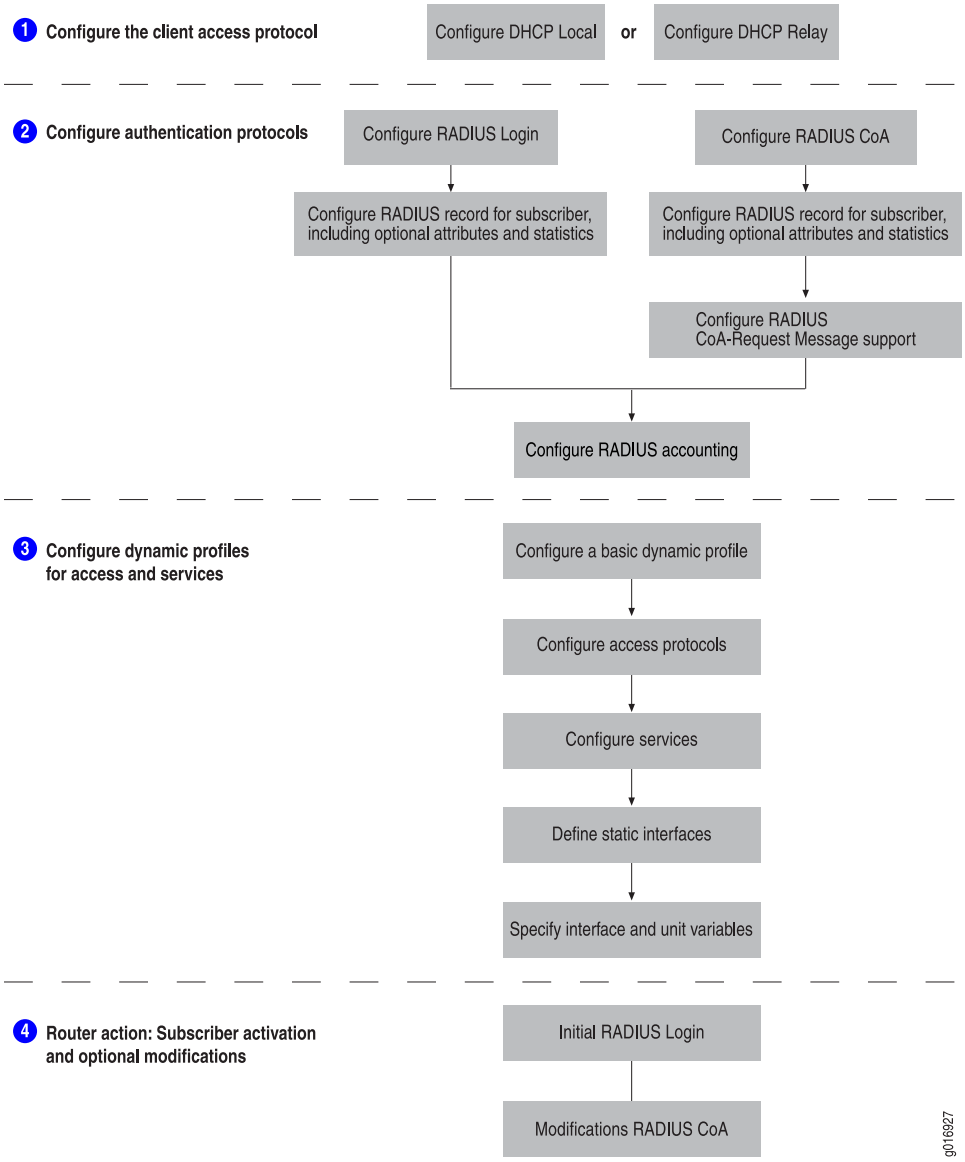

Related Documentation

- Subscriber Access Overview on page 3

Copyright © 2013, Juniper Networks, Inc.

- Subscriber Access Support Considerations on page 5
- Default Subscriber Service Overview on page 1127
- CLI-Activated Subscriber Services on page 1121

Collecting Subscriber Access Logs Before Contacting Juniper Technical Support

Problem When you experience a subscriber access problem in your network, we recommend that you collect certain logs before you contact Juniper 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 Technical Support in your request for assistance.

Solution To collect standard troubleshooting information:

- Redirect the command output to a file.

  user@host> **request support information | save rsi-1**

To configure logging to assist Juniper 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 © 2013, 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 Technical Support on page 15

Compressing Troubleshooting Logs from /var/logs to Send to Juniper Technical Support

Problem

You have collected logs on your device and need to send them to Juniper 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.

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
re1.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 Technical Support on page 13

Copyright © 2013, Juniper Networks, Inc.

PART 2

AAA Service Framework for Subscriber Access

- Configuring the AAA Service Framework for Subscriber Access on page 19
- Configuring Address-Assignment Pools for Subscriber Access on page 167
- Configuring Domain Maps for Subscriber Access on page 179
- AAA and Remote Subscriber Access Configuration Examples on page 193

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 2

Configuring the AAA Service Framework for Subscriber Access

- AAA Service Framework Overview on page 21
- Configuring Router or Switch Interaction with RADIUS Servers on page 23
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- Specifying the Authentication and Accounting Methods for Subscriber Access on page 24
- RADIUS Acct-On and Acct-Off Messages on page 25
- RADIUS Accounting Statistics for Subscriber Access Overview on page 26
- Understanding RADIUS Accounting Duplicate Reporting on page 28
- Configuring Per-Subscriber Session Accounting on page 29
- Retaining Authentication and Accounting Information During Session Startup on page 31
- Preservation of RADIUS Accounting Information During an Accounting Server Outage on page 32
- Configuring Per-Service Session Accounting on page 35
- Configuring RADIUS Server Parameters for Subscriber Access on page 36
- Specifying RADIUS Authentication and Accounting Servers for Subscriber Access on page 37
- Configuring Back-up Options for RADIUS Accounting on page 38
- Forcing the Router to Contact the Accounting Server Immediately on page 39
- Monitoring Pending RADIUS Accounting Stop Messages on page 39
- RADIUS Server Options for Subscriber Access on page 41
- Global RADIUS Options for Subscriber Access on page 44
- DNS Address Assignment Precedence on page 44
- Manual Configuration of the NAS-Port-ID RADIUS Attribute on page 44
- Manual Configuration of the NAS-Port-Type RADIUS Attribute on page 45
- Configuring RADIUS Server Options for Subscriber Access on page 47
- Configuring RADIUS Options for Subscriber Access Globally on page 50

Copyright © 2013, Juniper Networks, Inc.

- Subscriber Access Interface Description Storage and Reporting Through RADIUS Overview on page 50
- Configuring a NAS-Port-ID with Additional Options on page 54
- Configuring a Calling-Station-ID with Additional Attributes on page 56
- Configuring How RADIUS Attributes Are Used for Subscriber Access on page 58
- RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview on page 63
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65
- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 66
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 68
- Configuring the RADIUS NAS-Port-Type per VLAN on page 69
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 70
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 72
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 74
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 75
- Configuring the RADIUS NAS-Port Extended Format for ATM Interfaces on page 78
- Centrally Configured Opaque DHCP Options on page 80
- Monitoring DHCP Options Configured on RADIUS Servers on page 84
- Using RADIUS Dynamic Requests for Subscriber Access Management on page 86
- Dynamic Service Activation During Login Overview on page 86
- RADIUS-Initiated Change of Authorization (CoA) Overview on page 87
- RADIUS-Initiated Disconnect Overview on page 88
- Configuring RADIUS-Initiated Dynamic Request Support on page 90
- Verifying and Managing the RADIUS Dynamic-Request Feature on page 90
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 90
- RADIUS IETF Attributes Supported by the AAA Service Framework on page 91
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- DSL Forum Vendor-Specific Attributes on page 108
- AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 110
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 115
- DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS on page 119
- Error-Cause Codes (RADIUS Attribute 101) for Dynamic Requests on page 120
- Mapping Application Terminate Reasons and RADIUS Terminate Codes on page 121

 Copyright © 2013, Juniper Networks, Inc.

- AAA Terminate Reasons on page 123
- DHCP Terminate Reasons on page 124
- L2TP Terminate Reasons on page 124
- PPP Terminate Reasons on page 140
- Configuring Custom Terminate Reason Mappings on page 148
- Configuring an Access Profile for Subscriber Management on page 149
- Attaching Access Profiles on page 150
- DNS Name Server Address Overview on page 150
- Configuring DNS Name Server Addresses for Subscriber Management on page 151
- Understanding Session Options for Subscriber Access on page 153
- Configuring Subscriber Session Options on page 154
- Removing Inactive Dynamic Subscriber VLANs on page 155
- AAA Configuration Testing and Troubleshooting on page 156
- Testing a Subscriber AAA Configuration on page 156
- Tracing Subscriber Management Database Operations for Subscriber Access on page 158
- Configuring the Subscriber Management Database Trace Log Filename on page 159
- Configuring the Number and Size of Subscriber Management Database Log Files on page 159
- Configuring Access to the Subscriber Management Database Log File on page 160
- Configuring a Regular Expression for Subscriber Management Database Messages to Be Logged on page 160
- Configuring the Subscriber Management Database Tracing Flags on page 161
- Tracing Subscriber Management Session Database Replication Operations for Subscriber Access on page 161
- Configuring the Number and Size of Subscriber Management Session Database Replication Log Files on page 162
- Configuring a Regular Expression for Subscriber Management Session Database Replication Messages to Be Logged on page 163
- Configuring Access to the Subscriber Management Session Database Replication Log File on page 163
- Configuring the Subscriber Management Session Database Replication Trace Log Filename on page 164
- Configuring the Subscriber Management Session Database Replication Tracing Flags on page 164
- Verifying and Managing Subscriber AAA Information on page 164

AAA Service Framework Overview

The authentication, authorization, and accounting (AAA) Service Framework provides a single point of contact for all the authentication, authorization, accounting, address assignment, and dynamic request services that the router supports for network access.

The framework supports authentication and authorization through external servers, such as RADIUS. The framework also supports accounting and dynamic-request change of authorization (CoA) and disconnect operations through external servers, and address assignment through a combination of local address-assignment pools and RADIUS.

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

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

- Authentication—Authentication parameters defined in the access profile determine the authentication component of the AAA processing. For example, subscribers can be authenticated using an external authentication service such as RADIUS.
- Accounting—Accounting parameters in the access profile specify the accounting part of the AAA processing. For example, the parameters determine how the router collects and uses subscriber statistics. You can also configure AAA to enable the router to collect statistics on a per-service session basis for subscribers.
- RADIUS-initiated dynamic requests—A list of authentication server IP addresses in the access profile specify the RADIUS servers that can initiate dynamic requests to the router. Dynamic requests include CoA requests, which specify VSA modifications and service changes, and disconnect requests, which terminate subscriber sessions. The list of authentication servers also provide RADIUS-based dynamic service activation and deactivation during subscriber login.
- Address assignment—The AAA Service Framework assigns addresses to subscribers based on the configuration of local address-assignment pools. For example, the AAA framework collaborates with RADIUS servers to assign addresses from the specified pools.
- Subscriber secure policy—RADIUS VSAs and attributes provide RADIUS-initiated traffic mirroring on a per-subscriber basis.

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 23
- RADIUS Acct-On and Acct-Off Messages on page 25
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- Address-Assignment Pools Overview on page 167
- DNS Address Assignment Precedence on page 44
- RADIUS Accounting Statistics for Subscriber Access Overview on page 26
- Using RADIUS Dynamic Requests for Subscriber Access Management on page 86
- Subscriber Secure Policy Overview on page 1245

 Copyright © 2013, Juniper Networks, Inc.

Configuring Router or Switch Interaction with RADIUS Servers

You specify the RADIUS servers that the router or switch can use and you configure how the router or switch interacts with the servers. You can configure the router or switch to use multiple RADIUS servers on the network.

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

1. Configure the IP address of the RADIUS server and specify that you want to configure the router or switch interaction with the server.

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

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

   ```
   [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. (Optional) Configure the number of times that the router or switch attempts to contact a RADIUS accounting server. You can configure the router or switch to retry from 1 through 16 times. The default setting is 3 retry attempts.

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

5. 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 $nt1UE1*7688+
   ```

6. (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
   ```

7. 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 address configured on one of the router or switch interfaces.

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

Copyright © 2013, Juniper Networks, Inc.

8. (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 192.168.1.250]
user@host# set timeout 45
```

Related Documentation

- AAA Service Framework Overview on page 21
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 193

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 24.

2. Specify how accounting statistics are collected.

   See "Configuring Per-Subscriber Session Accounting" on page 29.

Related Documentation

- AAA Service Framework Overview on page 21
- Configuring Router or Switch Interaction with RADIUS Servers on page 23
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 193

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 © 2013, Juniper Networks, Inc.

You can specify the following authentication methods:

NOTE: For this release, 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. For this release, 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 23
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- Configuring Per-Subscriber Session Accounting on page 29
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 193

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.

Copyright © 2013, Juniper Networks, Inc.

- 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 21
- Configuring Per-Subscriber Session Accounting on page 29
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 115

RADIUS Accounting Statistics for Subscriber Access Overview

The AAA Service Framework enables you to configure how the router collects and uses accounting statistics for subscriber management.

For example, you can specify when statistics collection is terminated, the order in which different accounting methods are used, the types of statistics collected, and how often statistics are collected. You can also configure the router to request that the RADIUS server immediately update the accounting statistics when certain events occur, such as when a subscriber logs in or when a change of authorization (CoA) occurs.

Subscriber management provides two levels of subscriber accounting—subscriber session and service session. In subscriber session accounting, the router collects statistics for the entire subscriber session. In service session accounting, the router collects statistics for specific service sessions for the subscriber.

NOTE: **Subscriber management counts forwarded packets only. Dropped traffic (for example, as a result of a filter action) and control traffic are not included in the accounting statistics.**

The router uses the RADIUS attributes and Juniper Networks VSAs listed in Table 4 on page 27 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.

Copyright © 2013, Juniper Networks, Inc.

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 4: 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 29
- Configuring Per-Service Session Accounting on page 35
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 193

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 5 on page 28 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 5: 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 © 2013, Juniper Networks, Inc.

Table 5: 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.

Related Documentation

- Configuring Per-Subscriber Session Accounting on page 29

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 26 to provide the accounting statistics for the subscriber session.

To configure accounting for a subscriber session:

1. At the **[edit access profile *profile-name*]** hierarchy level, specify that you want to configure accounting.

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

2. (Optional) Configure AAA to issue an Acct-Stop message if the AAA server denies access to the subscriber.

```
[edit access profile profile-name accounting]
user@host# set accounting-stop-on-access-deny
```

3. (Optional) Configure AAA to send an Acct-Stop message if the subscriber fails AAA but is granted access by the AAA server.

```
[edit access profile profile-name accounting]
user@host# set accounting-stop-on-failure
```

4. (Optional) Configure the router or switch to send an Acct-Update message to the RADIUS accounting server when a CoA occurs.

```
[edit access profile profile-name accounting]
user@host# set coa-immediate-update
```

5. (Optional) Configure subscriber management to send the RADIUS accounting report to both the wholesaler and the retailer accounting servers.

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

6. (Optional) Configure the 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
```

7. (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
```

8. (Optional) Configure the order in which multiple accounting methods are used.

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

9. (Optional) Configure the types of statistics to gather. You can specify that the router or switch collect both volume and time statistics or only time statistics for subscriber sessions. When you change the type of statistics being collected, current subscribers continue to use the previous collection specification. Subscribers who log in after the change use the new specification.

```
[edit access profile profile-name accounting]
user@host# set statistics (time | volume-time)
```

10. (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

 Copyright © 2013, Juniper Networks, Inc.

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
```

11. (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
```

12. (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
```

13. (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 26
- Understanding RADIUS Accounting Duplicate Reporting on page 28
- Configuring Per-Service Session Accounting on page 35
- Retaining Authentication and Accounting Information During Session Startup on page 31
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 193

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 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 29

Preservation of RADIUS Accounting Information During an Accounting Server Outage

If the router loses contact with the RADIUS accounting server, as represented in Figure 4 on page 33, 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.

Copyright © 2013, Juniper Networks, Inc.

Figure 4: Topology with Loss of Access to Accounting Server

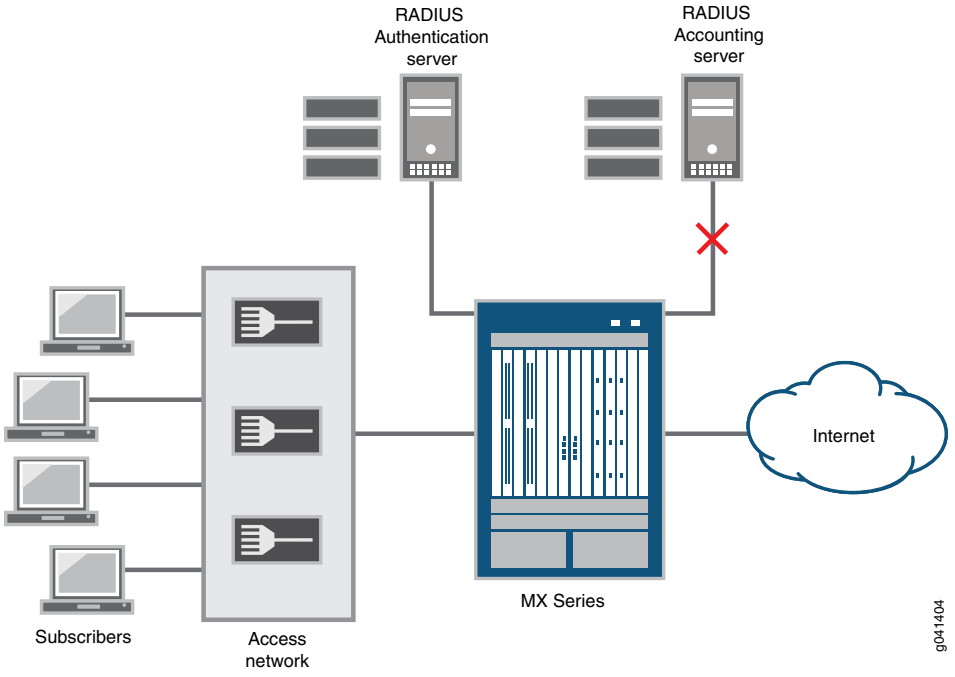

By default, the router must wait until the revert timer expires before it can attempt to contact the non-responsive server again. However, when you configure accounting backup, the revert timer is disabled and the router immediately retries its accounting requests as soon as the router fails to receive accounting acknowledgments. Accounting backup follows this sequence:

1. The router fails to receive accounting acknowledgments from the server.
2. The router immediately attempts to contact the accounting server and marks the server as offline if the router does not receive an acknowledgment before exhausting the number of retries.
3. The router next attempts to contact in turn each additional accounting server configured in the RADIUS profile.

   If a server is reached, then the router resumes sending accounting requests to this server.
4. If none of the servers responds or if no other servers are in the profile, the router declares a timeout and begins backing up the accounting data. It withholds all accounting stop messages and does not forward new accounting requests to the server.
5. During the outage, the router sends a single pending accounting stop message to the servers at periodic intervals.
6. If one of the servers acknowledges receipt, then the router sends all the pending stop messages to that server in batches at the same interval until all the stored stop messages have been sent. However, any new accounting requests are sent immediately rather being held and sent periodically.

Copyright © 2013, Juniper Networks, Inc.

The router replays accounting stop messages to the server in the correct order because it preserves both the temporal order among subscribers and the causal order between service and session stop requests for each subscriber. Only accounting stop messages are backed up, because they include the start time and duration of sessions and all the accounting statistics. This makes it unnecessary to withhold the accounting start messages, which eventually time out. Interim updates are not backed up and time out as well; if the session remains active, then the next interim update after the server connection is restored provides the interim accounting information.

You can configure the number of accounting stop messages that the router can queue pending restoration of contact with the accounting server. To preserve current accounting data in preference to collecting new accounting data, subscriber logins fail as soon as the maximum number of messages has been withheld. Subscriber logins resume immediately when the pending queue drops below the queue limit.

NOTE: **Service accounting stop messages are withheld for a maximum of ten services per subscriber. If a subscriber attempts to activate an eleventh service while that accounting server is offline, the activation fails.**

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

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

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

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

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

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

 Copyright © 2013, Juniper Networks, Inc.

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.

Related Documentation

- Configuring Back-up Options for RADIUS Accounting on page 38
- Forcing the Router to Contact the Accounting Server Immediately on page 39

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 26 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 6 on page 36 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 1175.

Table 6: 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 1175
- RADIUS Accounting Statistics for Subscriber Access Overview on page 26
- Configuring Per-Subscriber Session Accounting on page 29
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 193

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.

Copyright © 2013, Juniper Networks, Inc.

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 37.

3. Configure the RADIUS server options.

   See "Configuring RADIUS Server Options for Subscriber Access" on page 47.

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

   See "Configuring How RADIUS Attributes Are Used for Subscriber Access" on page 58.

Related Documentation

- Specifying RADIUS Authentication and Accounting Servers for Subscriber Access on page 37
- Configuring RADIUS Server Options for Subscriber Access on page 47
- Configuring How RADIUS Attributes Are Used for Subscriber Access on page 58

Specifying RADIUS Authentication and Accounting Servers for Subscriber Access

You can specify one or more RADIUS authentication or accounting servers to use for subscriber access management.

To configure RADIUS authentication and accounting support:

1. Specify that you want to configure RADIUS support.

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

2. Specify the IP address of the RADIUS server used for authentication.

   [edit access profile isp-bos-metro-fiber-basic radius]
   user@host# set authentication-server 192.168.1.251

3. Specify the IP address of the RADIUS server used for accounting.

   [edit access profile isp-bos-metro-fiber-basic radius]
   user@host# set accounting-server 192.168.1.250

To configure multiple RADIUS authentication or accounting servers:

- Specify the IP addresses of all RADIUS servers used for authentication or accounting.

  [edit access profile isp-bos-metro-fiber-basic radius]
  user@host# set authentication-server 192.168.1.251 192.168.1.252
  user@host# set accounting-server 192.168.1.250 192.168.1.251

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 23
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- Configuring RADIUS Server Options for Subscriber Access on page 47
- Configuring How RADIUS Attributes Are Used for Subscriber Access on page 58
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 193

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*

Copyright © 2013, Juniper Networks, Inc.

For example, the following statements configure the backup options for all subscriber accounting; these statements specify that the router holds no more than 32,000 pending accounting stops—at which point all subsequent subscriber logins fail—and holds them no longer than 6 hours—at which point all pending messages are flushed and subscriber logins resume if they were failing:

```
[edit access accounting-backup-options]
user@host# set max-pending-accounting-stops 32000
user@host# set max-withhold-time 360
```

Related Documentation

- Preservation of RADIUS Accounting Information During an Accounting Server Outage on page 32
- Forcing the Router to Contact the Accounting Server Immediately on page 39

Forcing the Router to Contact the Accounting Server Immediately

In the event of an accounting server outage while RADIUS accounting backup is enabled, by default the router waits for a time interval to expire before contacting the offline server. Rather than waiting for that interval to pass, you can force the router to immediately contact the server by issuing the **request network-access aaa replay pending-accounting-stops** command. The router sends a batch of pending accounting stop requests to the server. If the router receives an acknowledgment from the server, then the router continues to replay the pending messages to the server in batches at the periodic interval. If the router does not get that acknowledgment, then it resumes sending a single pending accounting stop message at the periodic interval.

To force the router to immediately contact the offline accounting server:

- Request the messages to be replayed.

  ```
  user@host> request network-access aaa replay pending-accounting-stops
  ```

Related Documentation

- Preservation of RADIUS Accounting Information During an Accounting Server Outage on page 32

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: 2010:9999:18::/48
Authentication State: AuthAcctStopAckWait
Accounting State: Acc-Stop-Stats-Pending
Service name: cos-service
  Service State: SvcInactive
  Session ID: 94
  Session uptime: 00:08:02
  Accounting status: on/time
  Service accounting session ID: 84:94-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
Service name: filter-service
  Service State: SvcInactive
  Session ID: 93
  Session uptime: 00:08:02
  Accounting status: on/volume+time
  Service accounting session ID: 84:93-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
Service name: filter-service6
  Service State: SvcInactive
  Session ID: 95
  Session uptime: 00:08:02
  Accounting status: on/volume+time
  Service accounting session ID: 84:95-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
```

You can display summary information for all users with a particular access profile. In the following example, only a single user, vjshah29@example.com, has the specified access profile, ce-ppp-profile:

```
user@host> show accounting pending-accounting-stops ce-ppp-profile

Type:         Username:               Session ID:     Service ID:       Service
  pppoe        vjshah29@example.com       84
  pppoe        vjshah29@example.com       84              94                cos-service
  pppoe        vjshah29@example.com       84              93                filter-service
  pppoe        vjshah29@example.com       84              95                filter-service6
```

You can also display summary information for all subscribers that have accounting-stop messages pending, regardless of access profile. The next example displays information for two users. Because the subscriber larry@example.com is not shown in the previous example, he must have a different access profile than vjshah29@example.com, even though he has received the same services.

```
user@host> show accounting pending-accounting-stops terse
```

 Copyright © 2013, Juniper Networks, Inc.

```
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 32
- Configuring Back-up Options for RADIUS Accounting on page 38

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.

The following list describes the RADIUS options you can configure:

- **accounting-session-id-format**—The format the router uses to identify the accounting session. The identifier can be in one of the following formats. The router uses **decimal** format by default.
  - **decimal**—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.

- **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.
- **access-loop-id-local**—The Agent-Remote-Id and Agent-Circuit-Id are 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.

- **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:

Copyright © 2013, Juniper Networks, Inc.

- **adapter-width** ***width***—Number of bits in the adapter field.
- **port-width** ***width***—Number of bits in the port 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 will fail.**

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 1 through 32
- **port-width**—Number of bits in the ATM port field, in the range 1 through 32
- **slot-width**—Number of bits in the ATM slot field, in the range 1 through 32
- **vpi-width**—Number of bits in the ATM virtual path identifier (VPI) field, in the range 1 through 32
- **vci-width**—Number of bits in the ATM virtual circuit identifier (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.**

- **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.
- **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 47

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 500 through 4000 requests per second. The default is 500 requests per second.

Related Documentation

- Configuring RADIUS Options for Subscriber Access Globally on page 50
- **request-rate on page 1943**
- **revert-interval on page 1947**

DNS Address Assignment Precedence

Subscriber management supports three 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 in the local address pool on the router.

Related Documentation

- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- Address-Assignment Pools Overview on page 167

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.

 Copyright © 2013, Juniper Networks, Inc.

You can configure access profiles to specify additional information in the NAS-Port-ID. The additional information can be any combination of the interface description (the default value), the Agent Circuit ID, the Agent Remote ID, and the NAS identifier. You can also specify an optional delimiter character, which separates the values in a NAS-Port-ID. The default delimiter character is the hash character (**#**).

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 54
- RADIUS Server Options for Subscriber Access on page 41

Manual Configuration of the NAS-Port-Type RADIUS Attribute

Subscriber management uses the NAS-Port-Type (RADIUS attribute 61) to identify the type of physical port that is used to authenticate subscribers. By default, subscriber management uses a NAS-Port-Type of **ethernet**.

You can optionally configure access profiles to provide the value for the NAS-Port-Type attribute, which enables you to explicitly specify the NAS port type that is used for a given connection. For example, you might configure an access profile that specifies that a NAS port type of **wireless** is used for all Ethernet connections that are managed by that access profile.

Copyright © 2013, Juniper Networks, Inc.

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 7 on page 46 shows the supported port type values for RADIUS attribute 61 (NAS-Port-Type) that you can include in an access profile.

Table 7: 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
piafs	6	Personal Handyphone System (PHS) Internet Access Forum Standard
sdsl	11	Symmetric DSL

Copyright © 2013, Juniper Networks, Inc.

Table 7: RADIUS NAS-Port-Type Values *(continued)*

Statement Option	NAS-Port-Type Value	Description
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 47

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.

   [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 stacked-vlan-width 10 vlan-width 2

   - 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).

Copyright © 2013, Juniper Networks, Inc.

```
[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).

```
[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
```

Related Documentation

- RADIUS Server Options for Subscriber Access on page 41
- Configuring Router or Switch Interaction with RADIUS Servers on page 23
- Manual Configuration of the NAS-Port-Type RADIUS Attribute on page 45
- Configuring a NAS-Port-ID with Additional Options on page 54
- Configuring a Calling-Station-ID with Additional Attributes on page 56

- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 193

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 **1000**

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 **86400**

Related Documentation

- Global RADIUS Options for Subscriber Access on page 44
- request-rate on page 1943
- revert-interval on page 1947

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 © 2013, 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 radius;
                    username-include {
                        user-prefix rich;
                    }
                }
                }
                interface ge-1/0/0.100;
            }
        }
    }
}
interfaces {
    ge-1/0/0 {
        description subscriber_ifd_descr;
        vlan-tagging;
        unit 100 {
            description subscriber_ifl_descr;
            vlan-id 100;
            family inet {
                unnumbered-address lo0.0 preferred-source-address 100.20.0.1;
```

```
            }
        }
    }
}
```

Reporting Interface Descriptions on Underlying Logical Interfaces

Underlying logical interfaces can apply to both DHCP and PPP.

For DHCP, Junos OS selects which interface description to report based on the following order of precedence:

1. Underlying logical interface description
2. Underlying physical interface description

NOTE: **For DHCP, Junos OS does not report the IP demux logical interface description.**

For PPP over an underlying VLAN or VLAN demux interface, Junos OS selects which interface description to report based on the following order of precedence:

1. PPP interface description
2. Underlying VLAN without a demux interface or VLAN demux logical interface description
3. Underlying physical interface description

Example: PPP over an Underlying VLAN Demux Interface

The following example shows a PPP subscriber over an underlying VLAN demux interface. This configuration includes three possible interface descriptions. Based on the order of precedence that Junos OS uses to select the interface description for PPP, the interface description is reported as subscriber_ppp_ifl_descr_0.

```
interfaces {
    ge-1/0/0 {
        description subscriber_ifd_descr;
        hierarchical-scheduler maximum-hierarchy-levels 2;
        flexible-vlan-tagging;
    }
    demux0 {
        unit 0 {
            vlan-tags outer 1 inner 1;
            description subscriber_under_ifl_descr_1_1;
            demux-options {
                underlying-interface ge-1/0/0;
            }
            family pppoe {
                duplicate-protection;
            }
        }
        unit 1 {
            vlan-tags outer 1 inner 2;
            description subscriber_under_ifl_descr_1_2;
```


Copyright © 2013, Juniper Networks, Inc.

```
            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 {
            dynamic-profile vlan-prof {
                accept inet;
                ranges {
```

```
                    any;
                }
            }
            authentication {
                password radius;
                username-include {
                    user-prefix rich;
                }
            }
        }
    }
}
```

Related Documentation

- report-interface-descriptions (Edit Access) on page 1943
- exclude on page 1625
- Configuring RADIUS Server Options for Subscriber Access on page 47
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24

Configuring a NAS-Port-ID with Additional Options

You can include optional values in the NAS-Port-ID (RADIUS attribute 87), which identifies the physical interface 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 values in the NAS-Port-ID:

- **agent-circuit-id**
- **agent-remote-id**
- **interface-description**
- **nas-identifier**

To configure an access profile to provide additional options in the NAS-Port-ID:

1. Specify the access profile you want to configure.

   [edit]
   user@host# **edit access profile retailer25**

2. Specify that you want to configure RADIUS options.

   [edit access profile retailer25]
   user@host# **edit radius options**

3. Specify the character to use as the delimiter between the different attribute values in the NAS-Port-ID. By default, subscriber management uses the hash character (**#**).

   [edit access profile retailer25 radius options]
   user@host# **set nas-port-delimiter %**

4. Specify that you want to configure the format of the NAS-Port-ID.

   [edit access profile retailer25 radius options]
   user@host# **edit** nas-port-id-format

Copyright © 2013, Juniper Networks, Inc.

5. Include the interface description in the NAS-Port-ID. (The interface description is not included by default when you configure the **nas-port-id-format** statement.)

   [edit access profile retailer25 radius options nas-port-id-format]
   user@host# **set interface-description**

6. Include the Agent Circuit ID in the NAS-Port-ID.

   [edit access profile retailer25 radius options nas-port-id-format]
   user@host# **set agent-circuit-id**

7. Include the Agent Remote ID in the NAS-Port-ID.

   [edit access profile retailer25 radius options nas-port-id-format]
   user@host# **set agent-remote-id**

8. Include the NAS identifier value in the NAS-Port-ID.

   [edit access profile retailer25 radius options nas-port-id-format]
   user@host# **set nas-identifier**

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 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**)—String that uniquely identifies 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**)—String that identifies the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request. The agent remote identifier (ARI) string is stored in 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**)—Description of the interface, which is not included in the Calling-Station-ID by default.
- NAS identifier (**nas-identifier**)—Name of the NAS that originated the authentication or accounting request. NAS-Identifier is RADIUS IETF attribute 32.

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.

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*]
   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*

Copyright © 2013, Juniper Networks, Inc.

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. Include the interface description in the Calling-Station-ID.

   [edit access profile *profile-name* radius options calling-station-id-format]
   user@host# **set interface-description**

7. 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**

8. 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**

9. 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**

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**.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 47
- RADIUS Server Options for Subscriber Access on page 41

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 8 on page 58 for the attributes you can configure.

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

4. Configure RADIUS to exclude the specified attribute from the specified RADIUS message type. See Table 9 on page 59 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 8 on page 58 lists the attributes supported in the **ignore** statement.

Table 8: 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

Copyright © 2013, Juniper Networks, Inc.

Table 8: Attributes That Can Be Ignored in RADIUS Access-Accept Messages *(continued)*

CLI Entry	Attribute Name	Attribute Number
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, Acct-On, Acct-Off, Acct-Start, and Acct-Stop messages. Table 9 on page 59 lists the attributes and message types supported in the **exclude** statement.

Table 9: 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
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

Copyright © 2013, Juniper Networks, Inc.

Table 9: Attributes That Can Be Excluded from RADIUS Messages *(continued)*

CLI Entry	Attribute Name	Attribute Number	Supported Message Type
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 9: Attributes That Can Be Excluded from RADIUS Messages *(continued)*

CLI Entry	Attribute Name	Attribute Number	Supported Message Type
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 9: Attributes That Can Be Excluded from RADIUS Messages *(continued)*

CLI Entry	Attribute Name	Attribute Number	Supported Message Type
interface-description	Interface-Desc	Juniper Networks VSA 26–53	Access-Request Accounting-Start Accounting-Stop
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
ouput-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
output-ipv6-packets	IPv6-Acct-Output-Packets	Juniper Networks VSA 26–154	Accounting-Stop

Copyright © 2013, Juniper Networks, Inc.

Table 9: Attributes That Can Be Excluded from RADIUS Messages *(continued)*

CLI Entry	Attribute Name	Attribute Number	Supported Message Type
upstream-calculated-qos-rate	Upstream-Calculated-QoS-Rate	Juniper Networks VSA 26-142	Access-Request Accounting-Start Accounting-Stop Interim-accounting

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 23
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- Specifying RADIUS Authentication and Accounting Servers for Subscriber Access on page 37
- Configuring RADIUS Server Options for Subscriber Access on page 47
- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 193

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 63
- NAS-Port RADIUS Attribute on page 64
- NAS-Port Options Configuration and Subscriber Network Access Models on page 64
- NAS-Port Options Definition on page 64

NAS-Port-Type RADIUS Attribute

The NAS-Port-Type attribute specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber. When you use the **nas-port-type** statement to configure the NAS-Port-Type, you can specify one of several predefined port types, or a user-defined port type value in the range 0 through 65535.

NAS-Port RADIUS Attribute

The NAS-Port attribute specifies the physical port number of the NAS that is authenticating the user, and is formed by a combination of the physical port's slot number, port number, adapter number, VLAN ID, and S-VLAN ID. The NAS-Port extended format, which you configure with the **nas-port-extended-format** statement, specifies the number of bits (bit width) for each field in the NAS-Port attribute: slot, adapter, port, VLAN, and S-VLAN.

To include stacked VLAN IDs, in addition to VLAN IDs, in the NAS-Port extended format, use the **stacked** option as part of the **nas-port-extended-format** statement. If you do not configure the **stacked** option, stacked VLAN IDs are not included in the extended format.

NAS-Port Options Configuration and Subscriber Network Access Models

Configuring the NAS-Port-Type and the extended format for NAS-Port on a per-VLAN, per-stacked VLAN, or per-physical interface basis is useful in network configurations that use the following subscriber access models:

- 1:1 access model (per-VLAN basis)—In a 1:1 access model, dedicated customer VLANs (C-VLANs) provide a one-to-one correspondence between an individual subscriber and the VLAN encapsulation.
- N:1 access model (per–S-VLAN basis)—In an N:1 access model, service VLANs are dedicated to a particular service, such as video, voice, or data, instead of to a particular subscriber. Because a service VLAN is typically shared by many subscribers within the same household or in different households, the N:1 access model provides a many-to-one correspondence between individual subscribers and the VLAN encapsulation.
- 1:1 or N:1 access model (per-physical interface basis)—You can configure the NAS-Port-Type and NAS-Port format on a per-physical interface basis for both the 1:1 access model and the N:1 access model.

NAS-Port Options Definition

As an alternative to globally configuring the NAS-Port-Type and NAS-Port extended format in an access profile, you can configure these attributes on a per-interface, per-VLAN, or per-stacked VLAN basis. To do so, you must create a *NAS-Port options definition*, which includes some or all of the following components:

- NAS-Port-Type value—Specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber.
- NAS-Port extended format—Configures the number of bits (bit width) for each field in the NAS-Port attribute, including: slot, adapter, port, VLAN, and S-VLAN. Optionally, you can also use the **stacked** option as part of the **nas-port-extended-format** statement to include S-VLAN IDs, in addition to VLAN IDs, in the extended format. If you do not configure the **stacked** option, stacked VLAN IDs are not included in the extended format.
- VLAN ranges or S-VLAN ranges—Defines the VLAN range of subscribers or stacked VLAN range of subscribers to which each NAS-Port options definition applies.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65
- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 66

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 63
- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 66

Copyright © 2013, 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 © 2013, 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 68.
   - For per-VLAN configurations, see "Configuring the RADIUS NAS-Port-Type per VLAN" on page 69.
   - For per-stacked VLAN configurations, see "Configuring the RADIUS NAS-Port-Type per Stacked VLAN" on page 70.
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 72.
   - For per-VLAN configurations, see "Configuring the RADIUS NAS-Port Extended Format per VLAN" on page 74.
   - For per-stacked VLAN configurations, see "Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN" on page 75.

Related Documentation

- RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview on page 63
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 68
- Configuring the RADIUS NAS-Port-Type per VLAN on page 69
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 70
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 72
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 74
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 75

Copyright © 2013, Juniper Networks, Inc.

Configuring the RADIUS NAS-Port-Type per Physical Interface

As an alternative to globally configuring the NAS-Port-Type (61) RADIUS attribute in an access profile, you can configure the NAS-Port-Type on a per-physical interface basis as part of a NAS-Port options definition. The NAS-Port-Type specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber.

Configuring NAS-Port options definitions on a per-physical interface basis is useful in network configurations that use a 1:1 access model or an N:1 access model.

To configure the NAS-Port-Type RADIUS attribute per physical interface:

1. Specify the interface you want to configure.

   [edit]
   user@host# **edit interfaces** ***interface-name***

2. Enable VLAN tagging on the interface.

   [edit interfaces *interface-name*]
   user@host# **set vlan-tagging**

   Setting VLAN tagging enables the reception and transmission of 802.1Q VLAN-tagged frames on the interface. You must enable VLAN tagging before you can configure the VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a physical interface.

   [edit interfaces *interface-name*]
   user@host# **edit** radius-options

4. Create a named NAS-Port options definition.

   [edit interfaces *interface-name* radius-options]
   user@host# **edit** nas-port-options ***nas-port-options-name***

5. Configure the NAS-Port-Type.

   [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]
   user@host# **set** nas-port-type ***port-type***

6. Configure the VLAN range or ranges to which the NAS-Port options definition applies.

   [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]
   user@host# **set** vlan-ranges (**any** | ***low-tag–high-tag***)

   Per-physical interface configurations typically require you to create a VLAN range that consists of all VLAN IDs on the physical interface. To do so, use the **any** option with the **vlan-ranges** statement.

The following example shows a per-interface NAS-Port options definition named subscribers-east that configures the **wireless-umts** NAS-Port-Type for a VLAN range consisting of all VLAN IDs on Gigabit Ethernet physical interface ge-1/0/0.

```
[edit interfaces ge-1/0/0 radius-options]
nas-port-options subscribers-east {
   nas-port-type wireless-umts;
   vlan-ranges {
```


Copyright © 2013, Juniper Networks, Inc.

```
        any;
    }
}
```

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65
- Configuring the RADIUS NAS-Port-Type per VLAN on page 69
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 70
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 72
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 74
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 75

Configuring the RADIUS NAS-Port-Type per VLAN

As an alternative to globally configuring the NAS-Port-Type (61) RADIUS attribute in an access profile, you can configure the NAS-Port-Type on a per-VLAN basis as part of a NAS-Port options definition. The NAS-Port-Type specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber.

Configuring NAS-Port options definitions on a per-VLAN basis is useful in network configurations that use a 1:1 access model.

To configure the NAS-Port-Type RADIUS attribute per VLAN:

1. Specify the interface you want to configure.

   [edit]
   user@host# **edit interfaces** *interface-name*

2. Enable VLAN tagging on the interface.

   [edit interfaces *interface-name*]
   user@host# **set vlan-tagging**

   Setting VLAN tagging enables the reception and transmission of 802.1Q VLAN-tagged frames on the interface. You must enable VLAN tagging before you can configure the VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a VLAN interface.

   [edit interfaces *interface-name*]
   user@host# **edit radius-options**

4. Create a named NAS-Port options definition.

   [edit interfaces *interface-name* radius-options]
   user@host# **edit nas-port-options** ***nas-port-options-name***

5. Configure the NAS-Port-Type.

   [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]

```
user@host# set nas-port-type port-type
```

6. Configure the VLAN range or ranges to which the NAS-Port options definition applies.

```
[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]
user@host# set vlan-ranges (any | low-tag–high-tag)
```

Per-VLAN configurations typically require you to create a VLAN range that consists of a single VLAN ID on the physical interface. To do so, set the ***low-tag*** and ***high-tag*** options in the **vlan-ranges** statement to the same value, as shown in the following example.

The following example shows a per-VLAN NAS-Port options definition named subscribers-west that configures the **ethernet** NAS-Port-Type for VLAN ID 3 on Gigabit Ethernet physical interface ge-1/1/0.

```
[edit interfaces ge-1/1/0 radius-options]
nas-port-options subscribers-west {
   nas-port-type ethernet;
   vlan-ranges {
      3-3;
   }
}
```

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 68
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 70
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 72
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 74
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 75

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
```

 Copyright © 2013, Juniper Networks, Inc.

2. Enable stacked VLAN tagging on the interface.

   [edit interfaces *interface-name*]
   user@host# **set stacked-vlan-tagging**

   Setting stacked VLAN tagging enables you to configure dual VLAN tags for all logical interfaces on the physical interface. You must enable stacked VLAN tagging before you can configure the stacked VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a stacked VLAN interface.

   [edit interfaces *interface-name*]
   user@host# **edit** radius-options

4. Create a named NAS-Port options definition.

   [edit interfaces *interface-name* radius-options]
   user@host# **edit** nas-port-options ***nas-port-options-name***

5. Configure the NAS-Port-Type.

   [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]
   user@host# **set** nas-port-type ***port-type***

6. Configure the stacked VLAN range or ranges to which the NAS-Port options definition applies.

   [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]
   user@host# **set** stacked-vlan-ranges (**any** | ***low-outer-tag–high-outer-tag***),**any**

   Per-stacked VLAN configurations require you to create a stacked VLAN range of subscribers to which the NAS-Port options definition applies. You must configure the low and high outer tags (VLAN IDs) in the range 1 through 4094, and the inner tag (S-VLAN ID) as **any** to represent all S-VLAN ID tags.

7. Repeat Steps 3 through 6 to configure additional NAS-Port options definitions on this interface.

The following example creates two NAS-Port options definitions, subscribers-north and subscribers-south, configured on a per-stacked VLAN basis on Gigabit Ethernet physical interface ge-1/1/0.

The subscribers-north definition configures a NAS-Port-Type user-defined value (4711) for a stacked VLAN range with outer VLAN ID 1 and all inner S-VLAN IDs. The subscribers-south definition configures a NAS-Port-Type user-defined value (4722) for a stacked VLAN range with outer VLAN IDs in the range 2 through 10, and all inner S-VLAN IDs.

```
[edit interfaces ge-2/0/1 radius-options]
nas-port-options subscribers-north {
   nas-port-type 4711;
   stacked-vlan-ranges {
      1-1,any;
   }
}
   nas-port-options subscribers-south {
      nas-port-type 4722;
```

```
    stacked-vlan-ranges {
      2-10,any;
    }
  }
```

Related Documentation

- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 68
- Configuring the RADIUS NAS-Port-Type per VLAN on page 69
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 72
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 74
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 75

Configuring the RADIUS NAS-Port Extended Format per Physical Interface

As an alternative to globally configuring the extended format for the NAS-Port (5) RADIUS attribute in an access profile, you can configure the NAS-Port extended format on a per-physical interface basis as part of a NAS-Port options definition. The NAS-Port extended format configures the number of bits (bit width) in each field in the NAS-Port attribute, including: slot, adapter, port, VLAN, and S-VLAN.

Configuring NAS-Port options definitions on a per-physical interface basis is useful in network configurations that use a 1:1 access model or an N:1 access model.

To configure an extended format for the NAS-Port RADIUS attribute per physical interface:

1. Specify the interface you want to configure.

   [edit]
   user@host# **edit interfaces** ***interface-name***

2. Enable VLAN tagging on the interface.

   [edit interfaces *interface-name*]
   user@host# **set vlan-tagging**

   Setting VLAN tagging enables the reception and transmission of 802.1Q VLAN-tagged frames on the interface. You must enable VLAN tagging before you can configure the VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a physical interface.

   [edit interfaces *interface-name*]
   user@host# **edit radius-options**

4. Create a named NAS-Port options definition.

   [edit interfaces *interface-name* radius-options]
   user@host# **edit nas-port-options** ***nas-port-options-name***

 Copyright © 2013, Juniper Networks, Inc.

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 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 68
- Configuring the RADIUS NAS-Port-Type per VLAN on page 69
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 70
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 74
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 75

Copyright © 2013, 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.

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.

 Copyright © 2013, Juniper Networks, Inc.

```
[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 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 68
- Configuring the RADIUS NAS-Port-Type per VLAN on page 69
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 70
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 72
- Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN on page 75

Configuring the RADIUS NAS-Port Extended Format per Stacked VLAN

As an alternative to globally configuring the extended format for the NAS-Port (5) RADIUS attribute in an access profile, you can configure the NAS-Port extended format on a per-stacked VLAN basis as part of a NAS-Port options definition. The NAS-Port extended format configures the number of bits (bit width) in each field in the NAS-Port attribute, including: slot, adapter, port, VLAN, and S-VLAN.

Configuring NAS-Port options definitions on a per-stacked VLAN basis is useful in network configurations that use an N:1 access model.

To configure an extended format for the NAS-Port RADIUS attribute per stacked VLAN:

1. Specify the interface you want to configure.

   [edit]
   user@host# **edit interfaces** ***interface-name***

2. Enable stacked VLAN tagging on the interface.

   [edit interfaces *interface-name*]
   user@host# **set stacked-vlan-tagging**

   Setting stacked VLAN tagging enables you to configure dual VLAN tags for all logical interfaces on the physical interface. You must enable stacked VLAN tagging before you can configure the stacked VLAN ranges to which the NAS-Port options definition applies.

3. Specify that you want to configure RADIUS options for a stacked VLAN interface.

   [edit interfaces *interface-name*]
   user@host# **edit** radius-options

4. Create a named NAS-Port options definition.

   [edit interfaces *interface-name* radius-options]
   user@host# **edit** nas-port-options ***nas-port-options-name***

5. Configure the NAS-Port extended format.

   [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]
   user@host# **set** nas-port-extended-format **slot-width** ***width*** **adapter-width** ***width*** **port-width** ***width*** **vlan-width** ***width*** **stacked**

   To include S-VLAN IDs, in addition to VLAN IDs, in the NAS-Port extended format, include the **stacked** option in the **nas-port-extended-format** statement.

6. Configure the stacked VLAN range or ranges to which the NAS-Port options definition applies.

   [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]
   user@host# **set** stacked-vlan-ranges (**any** | ***low-outer-tag—high-outer-tag***),**any**

   Per-stacked VLAN configurations require you to create a stacked VLAN range of subscribers to which the NAS-Port options definition applies. You must configure the low and high outer tags (VLAN IDs) in the range 1 through 4094, and the inner tag (S-VLAN ID) as **any** to represent all S-VLAN ID tags.

7. Repeat Steps 3 through 6 to configure additional NAS-Port options definitions on this interface.

The following example creates two NAS-Port options definitions, chicago-subscribers and barcelona-subscribers, configured on a per-stacked VLAN basis on Gigabit Ethernet physical interface ge-3/2/1.

The chicago-subscribers definition configures a NAS-Port extended format consisting of a 8-bit slot field, 8-bit adapter field, 8-bit port field, 4-bit stacked VLAN field, and 4-bit VLAN field. Because the **stacked** option is configured in this definition, S-VLAN IDs, in addition to VLAN IDs, are included in the extended format. The chicago-subscribers definition applies to a stacked VLAN range with outer VLAN ID 1, and all inner S-VLAN IDs.

The barcelona-subscribers definition configures a NAS-Port extended format consisting of a 8-bit slot field, 8-bit adapter field, 8-bit port field, 4-bit stacked VLAN field, and 4-bit VLAN field. Because the **stacked** option is *not* configured in this definition, S-VLAN IDs are not included in the extended format. The barcelona-subscribers definition applies to a stacked VLAN range with outer VLAN IDs in the range 2 through 10, and all inner S-VLAN IDs.

```
[edit interfaces ge-3/2/1 radius-options]
nas-port-options chicago-subscribers {
   nas-port-extended-format {
      slot-width 8;
      adapter-width 8;
      port-width 8;
```

 Copyright © 2013, Juniper Networks, Inc.

```
        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 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65
- Configuring the RADIUS NAS-Port-Type per Physical Interface on page 68
- Configuring the RADIUS NAS-Port-Type per VLAN on page 69
- Configuring the RADIUS NAS-Port-Type per Stacked VLAN on page 70
- Configuring the RADIUS NAS-Port Extended Format per Physical Interface on page 72
- Configuring the RADIUS NAS-Port Extended Format per VLAN on page 74

Configuring the RADIUS NAS-Port Extended Format for ATM Interfaces

As an alternative to globally configuring an extended format for the NAS-Port (5) RADIUS attribute in an access profile, you can configure the NAS-Port extended format on a per-physical interface basis for both Ethernet subscribers and ATM subscribers as part of a NAS-Port options definition. The NAS-Port extended format configures the number of bits (bit width) in each field of the NAS-Port attribute, including: slot, adapter, port, ATM virtual path identifier (VPI), and ATM virtual circuit identifier (VCI).

To configure the NAS-Port extended format for an ATM interface, include one or both of the following options in the **nas-port-extended-format** statement along with the other options as appropriate for your needs:

- **vpi-width**—Number of bits in the ATM VPI field, in the range 1 through 32
- **vci-width**—Number of bits in the ATM VCI field, in the range 1 through 32

NOTE: For ATM subscribers, the combined total of the widths of all fields must not exceed 32 bits, or the configuration fails. The router may truncate the values of individual fields depending on the bit width you specify.

To configure an extended format for the NAS-Port RADIUS attribute for an ATM interface:

1. Specify the ATM interface you want to configure.

   [edit]
   user@host# **edit interfaces** *interface-name*

2. Specify that you want to configure RADIUS options for a physical interface.

   [edit interfaces *interface-name*]
   user@host# **edit** radius-options

3. Create a named NAS-Port options definition.

   [edit interfaces *interface-name* radius-options]
   user@host# **edit** nas-port-options *nas-port-options-name*

4. Configure the NAS-Port extended format.

   [edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]
   user@host# **set** nas-port-extended-format **slot-width** *width* **adapter-width** *width* **port-width** *width* **vpi-width** *width* **vci-width** *width*

The following example shows a NAS-Port options definition named boston-subscribers for ATM interface at-1/0/4 that configures a NAS-Port extended format with an ATM slot width of 6 bits, ATM adapter width of 3 bits, ATM port width of 4 bits, ATM VPI width of 12 bits, and ATM VCI width of 24 bits.

```
[edit interfaces at-1/0/4 radius-options]
nas-port-options boston-subscribers {
   nas-port-extended-format {
      slot-width 6;
```


Copyright © 2013, Juniper Networks, Inc.

```
        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 66
- RADIUS Server Options for Subscriber Access on page 41
- Configuring RADIUS Server Options for Subscriber Access on page 47

Centrally Configured Opaque DHCP Options

Copyright © 2013, 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.**

- Data Flow for RADIUS-Sourced DHCP Options on page 82
- Multiple VSA 26-55 Instances Configuration on page 83
- DHCP Options That Cannot Be Centrally Configured on page 83

Data Flow for RADIUS-Sourced DHCP Options

Figure 5 on page 82 shows the procedure subscriber management (DHCP management) uses when configuring DHCP options for subscribers (DHCP clients).

Figure 5: DHCP Options Data Flow

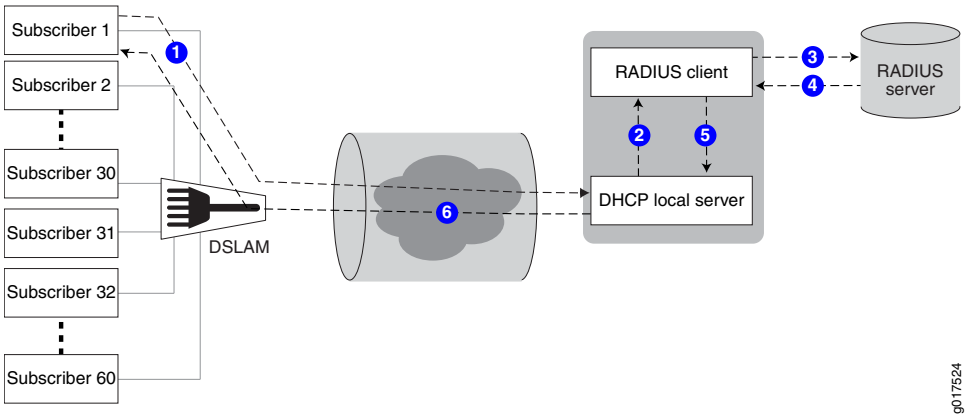

The following general sequence describes the data flow when subscriber management (DHCP management) uses RADIUS-sourced DHCP options and VSA 26-55 to configure a DHCP subscriber (client):

1. The subscriber (DHCP client) sends a DHCP discover message (or DHCPv6 solicit message) to the DHCP local server. The message includes client-sourced DHCP options.
2. The DHCP local server initiates authentication with the Junos OS RADIUS client.
3. The RADIUS client sends an Access-Request message on behalf of the subscriber (DHCP client) to the external RADIUS server. The message includes the subscriber's (DHCP client's) client-sourced DHCP options.
4. The external RADIUS server responds by sending an Access-Accept message to the RADIUS client. The Access-Accept message includes the RADIUS-sourced opaque DHCP options in VSA 26-55.
5. The RADIUS client sends the DHCP options string to DHCP local server. If there are multiple VSA 26-55 instances, the RADIUS client first assembles them into a single options string.
6. DHCP local server processes all options into the DHCP offer (or DHCPv6 reply) message, except for the RADIUS-sourced VSA 26-55 DHCP options. After processing all other options, DHCP local server then appends the unmodified VSA 26-55 DHCP options to the message and sends the message to the subscriber (DHCP client).

Copyright © 2013, Juniper Networks, Inc.

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 10 on page 84 shows the DHCP options that you must not centrally configure on the RADIUS server.

Table 10: 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 84

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: 00:10:95:00:00:98
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
```

 Copyright © 2013, Juniper Networks, Inc.

```
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
```

Meaning

```
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
```

The DHCP options output provides the following information:

- The **len** field is the total number of hex values in the message.
- The hex values specify the type, length, and value (TLV) of DHCP options, and are converted to decimal to identify the DHCP options, as defined in RFC 2132.

The number of hex values that make up a particular DHCP option varies, depending on the length of the option. For example, the first DHCP option specified in the output includes three sets of hex values (**35 01 01**). The first hex value (**35**) identifies the option type, the second value (**01**) indicates the length of the value entry, which is this case is one set of hex values. The third hex value (**01**) specifies the value for the DHCP option.

In the second DHCP option specification (**39 02 02 40**), the hex value **39** is the type, and the length of **02** specifies that two sets of hex entries make up the value for the option. Therefore, this option specification uses four sets of hex entries; one for the type (**39**), one to specify the length (**02**), and two for the option value (**02 40**).

The third DHCP option is specified by the hex values **3d 07 01 00 10 94 00 00 08**. The hex value **3d** is the type, followed by the length (**07**), which specifies that the next seven sets of hex entries make up the value for the option. Therefore, this option specification uses a total of nine sets of hex entries; one for the type (**3d**), one to specify the length (**07**), and seven for the value of the DHCP option (**01 00 10 94 00 00 08**).

Table 11 on page 85 describes the first two options in more detail.

Table 11: 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 80
- *show subscribers*

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.

Related Documentation

- Dynamic Service Activation During Login Overview on page 86
- RADIUS-Initiated Change of Authorization (CoA) Overview on page 87
- RADIUS-Initiated Disconnect Overview on page 88
- Configuring RADIUS-Initiated Dynamic Request Support on page 90
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 90
- Error-Cause Codes (RADIUS Attribute 101) for Dynamic Requests on page 120

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.

Related Documentation

- Using RADIUS Dynamic Requests for Subscriber Access Management on page 86

Copyright © 2013, Juniper Networks, Inc.

- Configuring RADIUS-Initiated Dynamic Request Support on page 90
- RADIUS-Initiated Disconnect Overview on page 88

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.

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)
- 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 12 on page 87 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 Acct-Session-ID, the attribute identifies the specific subscriber and session. 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.

Table 12: Identification Attributes

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

Table 13 on page 88 shows the session attributes for CoA operations. Any additional client attributes that you include depend on your particular session requirements.

Table 13: 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.

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 86
- Dynamic Service Activation During Login Overview on page 86
- RADIUS-Initiated Disconnect Overview on page 88
- Configuring RADIUS-Initiated Dynamic Request Support on page 90

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.

Copyright © 2013, Juniper Networks, Inc.

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.

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 86
- Dynamic Service Activation During Login Overview on page 86

- Configuring RADIUS-Initiated Dynamic Request Support on page 90

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

Related Documentation

- Using RADIUS Dynamic Requests for Subscriber Access Management on page 86
- Dynamic Service Activation During Login Overview on page 86
- RADIUS-Initiated Change of Authorization (CoA) Overview on page 87
- RADIUS-Initiated Disconnect Overview on page 88
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 90
- Error-Cause Codes (RADIUS Attribute 101) for Dynamic Requests on page 120

Verifying and Managing the RADIUS Dynamic-Request Feature

Purpose Display RADIUS dynamic request statistics and information.

Action

- To display RADIUS dynamic request statistics:

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

Related Documentation

- *Junos OS Operational Mode Commands*

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.

 Copyright © 2013, Juniper Networks, Inc.

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

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

Related Documentation

- RADIUS IETF Attributes Supported by the AAA Service Framework on page 91
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- DSL Forum Vendor-Specific Attributes on page 108
- AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 110
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 115
- Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs on page 660

RADIUS IETF Attributes Supported by the AAA Service Framework

Table 14 on page 91 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 14: Supported RADIUS IETF Attributes

Attribute Number	Attribute Name	Description	Dynamic CoA Support
1	User-Name	• Name of user to be authenticated. • Configurable username override.	No
2	User-Password	• Password of user to be authenticated by Password Authentication Protocol (PAP). • Configurable password override.	No
4	NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user.	No

Table 14: 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 © 2013, Juniper Networks, Inc.

Table 14: 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

Table 14: 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 © 2013, Juniper Networks, Inc.

Table 14: 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
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 14: Supported RADIUS IETF Attributes *(continued)*

Attribute Number	Attribute Name	Description	Dynamic CoA Support
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 © 2013, Juniper Networks, Inc.

Table 14: 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:172.20.0.1:192.168.0.2: 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 © 2013, Juniper Networks, Inc.

Table 14: 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 110
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 115

Juniper Networks VSAs Supported by the AAA Service Framework

Table 15 on page 99 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 © 2013, Juniper Networks, Inc.

Table 15: 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]. The standard attribute is recommended 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

Table 15: 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
26-25	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	The 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 © 2013, Juniper Networks, Inc.

Table 15: Supported Juniper Networks VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-51	Tunnel-Disconnect-Cause-Info	Disconnect cause when a tunneled subscriber is disconnected, and the termination is initiated by the L2TP layer of the LNS. 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: *tunnel-disconnect-cause-info*	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

Table 15: 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*	No
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
26-70	Ignore-DF-Bit	State of the Ignore Don't Fragment (DF) bit on client interface	integer: • 0 = do not ignore • 1 = ignore	No
26-71	IGMP-Access-Name	Access list to use for the group (G) filter.	string: 32-octet	Yes
26-72	IGMP-Access-Src-Name	Access list to use for the source-group (S,G) filter.	string: 32-octet	Yes

Copyright © 2013, Juniper Networks, Inc.

Table 15: Supported Juniper Networks VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
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)	

Copyright © 2013, Juniper Networks, Inc.

Table 15: 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)	
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
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	
26–111	Acc-Aggr-Cir-Id-Bin	Unique identification of the DSL line.	integer: 8-octet	

Copyright © 2013, Juniper Networks, Inc.

Table 15: Supported Juniper Networks VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26–112	Acc-Aggr-Cir-Id-Asc	Identification of the uplink on the access node, as 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	
26–113	Act-Data-Rate-Up	Actual upstream data rate of the subscriber's synchronized DSL link.	integer: 4-octet	
26–114	Act-Data-Rate-Dn	Actual downstream data rate of the subscriber's synchronized DSL link.	integer: 4-octet	
26–115	Min-Data-Rate-Up	Minimum upstream data rate configured for the subscriber.	integer: 4-octet	
26–116	Min-Data-Rate-Dn	Minimum downstream data rate configured for the subscriber.	integer: 4-octet	
26–117	Att-Data-Rate-Up	Maximum upstream data rate that the subscriber can attain.	integer: 4-octet	
26–118	Att-Data-Rate-Dn	Maximum downstream data rate that the subscriber can attain.	integer: 4-octet	
26–119	Max-Data-Rate-Up	Maximum upstream data rate configured for the subscriber.	integer: 4-octet	
26–120	Max-Data-Rate-Dn	Maximum downstream data rate configured for the subscriber.	integer: 4-octet	
26–121	Min-LP-Data-Rate-Up	Minimum upstream data rate in low power state configured for the subscriber.	integer: 4-octet	
26–122	Min-LP-Data-Rate-Dn	Minimum downstream data rate in low power state configured for the subscriber.	integer: 4-octet	
26–123	Max-Interlv-Delay-Up	Maximum one-way upstream interleaving delay configured for the subscriber.	integer: 4-octet	
26–124	Act-Interlv-Delay-Up	Subscriber's actual one-way upstream interleaving delay..	integer: 4-octet	

Copyright © 2013, Juniper Networks, Inc.

Table 15: Supported Juniper Networks VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26–125	Max-Interlv-Delay-Dn	Maximum one-way downstream interleaving delay configured for the subscriber.	integer: 4-octet	
26–126	Act-Interlv-Delay-Dn	Subscriber's actual one-way downstream interleaving delay.	integer: 4-octet	
26–127	DSL-Line-State	State of the DSL line.	integer: 4-octet • 1 = Show uptime • 2 = Idle • 3 = Silent	
26–128	DSL-Type	Encapsulation used by the subscriber associated with the DSLAM interface from which requests are initiated.		
26-130	Qos-Set-Name	Interface set to apply to the dynamic profile.	string: *interface-set-name*	No
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	
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	
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

Copyright © 2013, Juniper Networks, Inc.

Table 15: Supported Juniper Networks VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26-146	CoS-Scheduler-Pmt-Type	CoS scheduler parameter type and description: • 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
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
26-161	IPv6-Delegated-Pool-Name	Address pool used to locally allocate a delegated prefix (IA_PD).	string	No
26-162	Tx-Connect-Speed	Indication of user's connection.	string	No
26-163	Rx-Connect-Speed	Indication of 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 • 2 = op-script	No
26-174	Client-Profile-Name	Enables the RADIUS server to override the client dynamic profile in the Access-Accept message.	string	No

Copyright © 2013, Juniper Networks, Inc.

Table 15: Supported Juniper Networks VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value	Dynamic CoA Support
26–177	Cos-Shaping-Rate	Effective downstream shaping rate for subscriber.	string	No

Related Documentation

- AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 110
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 115

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 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 16 on page 108 describes the DSL Forum VSAs.

Table 16: DSL Forum VSAs

Attribute Number	Attribute Name	Description	Value
[26-1]	Agent-Circuit-Id	Identifier for the subscriber agent circuit ID that corresponds to the DSLAM interface from which subscriber requests are initiated	string
[26-2]	Agent-Remote-Id	Unique identifier for the subscriber associated with the DSLAM interface from which requests are initiated	string
[26-129]	Actual-Data-Rate-Upstream	Actual upstream data rate of the subscriber's synchronized DSL link	integer: 4-octet

Copyright © 2013, Juniper Networks, Inc.

Table 16: DSL Forum VSAs *(continued)*

Attribute Number	Attribute Name	Description	Value
[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
[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

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS on page 119

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

Table 17 on page 110 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 17: 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	–	✓	–	–	–	–
26-5	Secondary-DNS	–	✓	–	–	–	–
26-6	Primary-WINS	–	✓	–	–	–	–

Copyright © 2013, Juniper Networks, Inc.

Table 17: AAA Access Messages—Supported RADIUS Attributes and Juniper Networks VSAs *(continued)*

Attribute Number	Attribute Name	Access Request	Access Accept	Access Reject	Access Challenge	CoA Request	Disconnect Request
26-7	Secondary-WINS	–	✓	–	–	–	–
26-8	Tunnel-Virtual-Router	–	✓	–	–	–	–
26-9	Tunnel-Password	–	✓	–	–	–	–
26-10	Ingress-Policy-Name	–	✓	–	–	–	–
26-11	Egress-Policy-Name	–	✓	–	–	–	–
26-23	IGMP-Enable	–	✓	–	–	–	–
26-25	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	✓	–	–	–	–	–
26-64	Tunnel-Group	–	✓	–	–	–	–
26-65	Activate-Service	–	✓	–	–	✓	–

Copyright © 2013, Juniper Networks, Inc.

Table 17: 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
26-66	Deactivate-Service	–	✓	–	–	✓	–
26-69	Service-Statistics	–	✓	–	–	✓	–
26-70	Ignore-DF-Bit	–	✓	–	–	–	–
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	–	✓	–	–	–	–
26-97	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	✓	–	–	–	–	–
26-113	Act-Data-Rate-Up	✓	–	–	–	–	–
26-114	Act-Data-Rate-Dn	✓	–	–	–	–	–
26-115	Min-Data-Rate-Up	✓	–	–	–	–	–
26-116	Min-Data-Rate-Dn	✓	–	–	–	–	–

Copyright © 2013, Juniper Networks, Inc.

Table 17: 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
26-117	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	–	✓	–	–	–	–
26-160	Vlan-Map-ID	–	✓	–	–	–	–
26-161	IPv6-Delegated-Pool-Name	–	✓	–	–	–	–
26-162	Tx-Connect-Speed	✓	–	–	–	–	–
26-163	Rx-Connect-Speed	✓	–	–	–	–	–

Copyright © 2013, Juniper Networks, Inc.

Table 17: 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
26-173	Service-Activate-Type	–	✓	–	–	✓	–
26-174	Client-Profile-Name	–	✓	–	–	–	–
27	Session-Timeout	–	✓	–	✓	–	–
31	Calling-Station-ID	✓	–	–	–	✓	–
32	NAS-Identifier	✓	–	–	–	–	–
44	Acct-Session-ID	✓	–	–	–	✓	✓
61	NAS-Port-Type	✓	–	–	–	–	–
64	Tunnel-Type	–	✓	–	–	–	–
65	Tunnel-Medium-Type	–	✓	–	–	–	–
66	Tunnel-Client-Endpoint	–	✓	–	–	–	–
67	Tunnel-Server-Endpoint	–	✓	–	–	–	–
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	–	✓	–	–	–	–

Copyright © 2013, Juniper Networks, Inc.

Table 17: 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
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 115
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 90
- RADIUS IETF Attributes Supported by the AAA Service Framework on page 91
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98

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

Table 18 on page 115 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 18: 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	–	✓	✓	–	–

Copyright © 2013, Juniper Networks, Inc.

Table 18: 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
22	Framed-Route	✓	✓	✓	–	–
25	Class	✓	✓	✓	–	–
26-10	Ingress-Policy-Name	✓	✓	✓	–	–
26-11	Egress-Policy-Name	✓	✓	✓	–	–
26-42	Input-Gigapackets	–	✓	✓	–	–
26-43	Output-Gigapackets	–	✓	✓	–	–
26-47	Ipv6-Primary-DNS	✓	✓	✓	–	–
26-48	Ipv6-Secondary-DNS	✓	✓	✓	–	–
26-51	Tunnel-Disconnect-Cause-Info	–	✓	–	–	–
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	✓	✓	✓	–	–

Copyright © 2013, Juniper Networks, Inc.

Table 18: 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
26-119	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	✓	✓	✓	–	–
26-141	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	–	✓	✓	–	–
26-162	Tx-Connect-Speed	✓	✓	✓	–	–
26-163	Rx-Connect-Speed	✓	✓	✓	–	–
26–177	Cos-Shaping-Rate	✓	✓	✓	–	–
31	Calling-Station-ID	✓	✓	✓	–	–
32	NAS-Identifier	✓	✓	✓	–	–

Copyright © 2013, Juniper Networks, Inc.

Table 18: 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
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	✓	✓	✓	–	–
82	Tunnel-Assignment-Id	✓	✓	✓	–	–
87	NAS-Port-Id	✓	✓	✓	–	–
90	Tunnel-Client-Auth-Id	✓	✓	✓	–	–
91	Tunnel-Server-Auth-Id	✓	✓	✓	–	–
99	Framed-IPv6-Route	✓	✓	✓	–	–

Copyright © 2013, Juniper Networks, Inc.

Table 18: 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
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 110
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 90
- RADIUS IETF Attributes Supported by the AAA Service Framework on page 91
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98

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

Table 19 on page 119 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 19: 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	✓	✓	✓	✓	✓
[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	✓	✓	✓	✓	–

Copyright © 2013, Juniper Networks, Inc.

Table 19: DSL Forum VSAs—Supported RADIUS Messages *(continued)*

Attribute Number	Attribute Name	Access Request	Acct Start	Acct Stop	Interim Acct	CoA Request
[26-137]	Minimum-Data-Rate-Upstream-Low-Power	✓	✓	✓	✓	–
[26-138]	Minimum-Data-Rate-Downstream-Low-Power	✓	✓	✓	✓	–
[26-139]	Maximum-Interleaving-Delay-Upstream	✓	✓	✓	✓	–
[26-140]	Actual-Interleaving-Delay-Upstream	✓	✓	✓	✓	–
[26-141]	Maximum-Interleaving-Delay-Downstream	✓	✓	✓	✓	–
[26-142]	Actual-Interleaving-Delay-Downstream	✓	✓	✓	✓	–
[26-144]	Access-Loop-Encapsulation	✓	✓	✓	✓	–
[26-254]	IWF-Session	✓	✓	✓	✓	–

Related Documentation

- DSL Forum Vendor-Specific Attributes on page 108

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 20 on page 120 describes the error-cause codes.

Table 20: 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.

Copyright © 2013, Juniper Networks, Inc.

Table 20: Error-Cause Codes (RADIUS Attribute 101) *(continued)*

Code	Value	Description
506	Resources unavailable	A request could not be honored due to lack of available NAS resources (such as memory).

Related Documentation

- RADIUS-Initiated Change of Authorization (CoA) Overview on page 87
- RADIUS-Initiated Disconnect Overview on page 88

Mapping Application Terminate Reasons and RADIUS Terminate Codes

The Junos OS software uses default configuration mapping of terminate reasons for various protocols (AAA, DHCP, L2TP, and PPP) to RADIUS Acct-Terminate-Cause attributes. You can optionally create customized mappings between a terminate reason and a RADIUS Acct-Terminate-Cause attribute, enabling you to provide different information about the cause of a termination.

When a AAA, DHCP, L2TP, or PPP session is terminated, the router logs a message for the internal terminate reason and logs another message for the RADIUS Acct-Terminate-Cause attribute (RADIUS attribute 49). RADIUS attribute 49 is also included in RADIUS Acct-Off and Acct-Stop messages. You can use the logged information to help monitor and troubleshoot terminated sessions.

NOTE: **A single mapping for RADIUS account termination is shared by all clients.**

Table 21 on page 121 lists the IETF RADIUS Acct-Terminate-Cause codes that you can use to map application terminate reasons. In addition, you can also configure and use proprietary codes for values beyond 22.

Table 21: Supported RADIUS Acct-Terminate-Cause Codes

Code	Name	Description
1	User Request	User initiated the disconnect (log out)
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

Copyright © 2013, Juniper Networks, Inc.

Table 21: Supported RADIUS Acct-Terminate-Cause Codes *(continued)*

Code	Name	Description
7	Admin Reboot	System administrator terminated the session on the NAS; for example, prior to rebooting the NAS
8	Port Error	NAS detected an error on the port that required ending the session
9	NAS Error	NAS detected an error (other than on the port) that required ending the session
10	NAS Request	NAS ended the session for a non-error reason
11	NAS Reboot	NAS ended the session due to a non-administrative reboot
12	Port Unneeded	NAS ended the session because the resource usage fell below the low threshold; for example, the bandwidth-on-demand algorithm determined that the port was no longer needed
13	Port Preempted	NAS ended the session to allocate the port to a higher-priority use
14	Port Suspended	NAS ended the session to suspend a virtual session
15	Service Unavailable	NAS was unable to provide the requested service
16	Callback	NAS is terminating the current session in order to perform callback for a new session
17	User Error	An error in the user input caused the session to be terminated
18	Host Request	The login host terminated the session normally
19	Supplicant Restart	Supplicant state machine was reinitialized
20	Reauthentication Failure	A previously authenticated supplicant failed to reauthenticate successfully following expiration of the reauthentication timer or explicit reauthentication request by management action
21	Port Reinitialized	The port's MAC has been reinitialized
22	Port Administratively Disabled	The port has been administratively disabled

Related Documentation

- Configuring Custom Terminate Reason Mappings on page 148
- AAA Terminate Reasons on page 123
- DHCP Terminate Reasons on page 124
- L2TP Terminate Reasons on page 124
- PPP Terminate Reasons on page 140

Copyright © 2013, Juniper Networks, Inc.

AAA Terminate Reasons

Table 22 on page 123 lists the default AAA terminate mappings. The table indicates the supported AAA deny and shutdown reasons and the RADIUS Acct-Terminate-Cause attributes they are mapped to by default.

Table 22: Default AAA Mappings

AAA Deny or Shutdown Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
deny address allocation failure	17	user error
deny no-resources	10	NAS request
deny server- request-timeout	17	user error
shutdown idle-timeout	4	idle timeout NOTE: For Junos release 13.1, the default mapping is to RADIUS accounting NAS Request (10) terminate cause.
shutdown session-timeout	5	session timeout NOTE: For Junos release 13.1, the default mapping is to RADIUS accounting NAS Request (10) terminate cause.
shutdown administrative-reset	6	admin reset
shutdown remote-reset	10	NAS request

NOTE: In Junos OS releases prior to releases 12.2R5, 12.3R3, and 13.2R1, the **idle-timeout** and **session-timeout** terminate reasons both had a default mapping to RADIUS accounting NAS Request (10) terminate cause. To support backward compatibility for these releases, you can configure the router to support the previous mapping by using the **terminate-code aaa shutdown (idle-timeout | session-timeout) radius 10** statement at the **[edit access]** hierarchy level.

Related Documentation

- Mapping Application Terminate Reasons and RADIUS Terminate Codes on page 121
- Configuring Custom Terminate Reason Mappings on page 148

Copyright © 2013, Juniper Networks, Inc.

DHCP Terminate Reasons

Table 23 on page 124 lists the default DHCP terminate mappings. The table indicates the supported DHCP terminate reasons and the RADIUS Acct-Terminate-Cause attributes they are mapped to by default.

Table 23: Default DHCP Mappings

DHCP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
nak	15	service unavailable
nas logout	10	NAS request
no offers	4	idle timeout
lost-carrier	2	session terminated / modem dropped DCD
client request	1	user request

Related Documentation

- Mapping Application Terminate Reasons and RADIUS Terminate Codes on page 121
- Configuring Custom Terminate Reason Mappings on page 148

L2TP Terminate Reasons

Table 24 on page 124 lists the default L2TP terminate mappings. The table indicates the supported L2TP terminate reasons and the RADIUS Acct-Terminate-Cause attributes they are mapped to by default.

Table 24: Default L2TP Mappings

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session create failed limit reached	9	NAS error
session create failed no resources	9	NAS error
session create failed single shot tunnel already fired	9	NAS error
session create failed too busy	9	NAS error
session failover protocol resync disconnect	6	admin reset
session hardware unavailable	8	port error
session no resources server port	9	NAS error
session not ready	9	NAS error
session rx cdn	10	NAS request
session rx cdn avp bad hidden	10	NAS request
session rx cdn avp bad value assigned session id	10	NAS request
session rx cdn avp duplicate value assigned session id	10	NAS request
session rx cdn avp malformed bad length	10	NAS request
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session rx iccn avp bad value unsupported proxy authen type	10	NAS request
session rx iccn avp malformed bad length	10	NAS request
session rx iccn avp malformed truncated	10	NAS request
session rx iccn avp missing mandatory connect speed	10	NAS request
session rx iccn avp missing mandatory framing type	10	NAS request
session rx iccn avp missing mandatory proxy authen challenge	10	NAS request
session rx iccn avp missing mandatory proxy authen id	10	NAS request
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session rx icrp no resources	10	NAS request
session rx icrp unexpected	10	NAS request
session rx icrq admin close	6	admin reset
session rx icrq authenticate failed host	10	NAS request
session rx icrq avp bad hidden	10	NAS request
session rx icrq avp bad value assigned session id	10	NAS request
session rx icrq avp bad value bearer type	10	NAS request
session rx icrq avp bad value cisco nas port	10	NAS request
session rx icrq avp duplicate value assigned session id	10	NAS request
session rx icrq avp malformed bad length	10	NAS request
session rx icrq avp malformed truncated	10	NAS request
session rx icrq avp missing mandatory assigned session id	10	NAS request
session rx icrq avp missing mandatory call serial number	10	NAS request
session rx icrq avp missing random vector	10	NAS request
session rx icrq avp missing secret	10	NAS request
session rx icrq avp unknown	10	NAS request
session rx icrq no resources	10	NAS request
session rx icrq unexpected	10	NAS request
session rx occn avp bad hidden	10	NAS request
session rx occn avp bad value framing type	10	NAS request
session rx occn avp malformed bad length	10	NAS request
session rx occn avp malformed truncated	10	NAS request
session rx occn avp missing mandatory connect speed	10	NAS request

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session rx sli avp unknown	10	NAS request
session rx sli no resources	10	NAS request
session rx unexpected packet lac incoming	10	NAS request
session rx unexpected packet lac outgoing	10	NAS request
session rx unexpected packet lns incoming	10	NAS request
session rx unexpected packet lns outgoing	10	NAS request
session rx unknown session id	10	NAS request
session rx wen avp bad hidden	10	NAS request
session rx wen avp malformed bad length	10	NAS request
session rx wen avp malformed truncated	10	NAS request
session rx wen avp missing mandatory call errors	10	NAS request
session rx wen avp missing random vector	10	NAS request
session rx wen avp missing secret	10	NAS request
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel host profile deleted	6	admin reset
tunnel rx scccn authenticate failed challenge	17	user error
tunnel rx scccn avp bad hidden	15	service unavailable
tunnel rx scccn avp bad value challenge response	15	service unavailable
tunnel rx scccn avp malformed bad length	15	service unavailable
tunnel rx scccn avp malformed truncated	15	service unavailable
tunnel rx scccn avp missing challenge response	17	user error
tunnel rx scccn avp missing random vector	15	service unavailable
tunnel rx scccn avp missing secret	15	service unavailable
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx sccrp avp bad value protocol version	15	service unavailable
tunnel rx sccrp avp bad value receive window size	15	service unavailable
tunnel rx sccrp avp duplicate value assigned tunnel id	15	service unavailable
tunnel rx sccrp avp malformed bad length	15	service unavailable
tunnel rx sccrp avp malformed truncated	15	service unavailable
tunnel rx sccrp avp missing challenge response	17	user error
tunnel rx sccrp avp missing mandatory assigned tunnel id	15	service unavailable
tunnel rx sccrp avp missing mandatory framing capabilities	15	service unavailable
tunnel rx sccrp avp missing mandatory host name	15	service unavailable
tunnel rx sccrp avp missing mandatory protocol version	15	service unavailable
tunnel rx sccrp avp missing random vector	15	service unavailable
tunnel rx sccrp avp missing secret	15	service unavailable
tunnel rx sccrp avp unexpected challenge response	15	service unavailable
tunnel rx sccrp avp unexpected challenge without secret	15	service unavailable
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx sccrq avp bad value challenge	15	service unavailable
tunnel rx sccrq avp bad value failover capability	15	service unavailable
tunnel rx sccrq avp bad value framing capabilities	15	service unavailable
tunnel rx sccrq avp bad value protocol version	15	service unavailable
tunnel rx sccrq avp bad value receive window size	15	service unavailable
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
tunnel rx stopccn	1	user request
tunnel rx stopccn avp bad hidden	15	service unavailable

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx recovery sccrp avp unexpected challenge response	15	service unavailable
tunnel rx recovery sccrp avp unexpected challenge without secret	15	service unavailable
tunnel rx recovery sccrp avp unknown	15	service unavailable
tunnel rx recovery sccrp no resources	15	service unavailable
tunnel rx recovery sccrp session id not null	15	service unavailable
tunnel rx recovery sccrq admin close	6	admin reset
tunnel rx recovery sccrq avp bad hidden	15	service unavailable
tunnel rx recovery sccrq avp bad value assigned tunnel id	15	service unavailable
tunnel rx recovery sccrq avp bad value bearer capabilities	15	service unavailable
tunnel rx recovery sccrq avp bad value challenge	15	service unavailable
tunnel rx recovery sccrq avp bad value framing capabilities	15	service unavailable
tunnel rx recovery sccrq avp bad value protocol version	15	service unavailable
tunnel rx recovery sccrq avp bad value receive window size	15	service unavailable
tunnel rx recovery sccrq avp bad value tunnel recovery	15	service unavailable
tunnel rx recovery sccrq avp duplicate value assigned tunnel id	15	service unavailable
tunnel rx recovery sccrq avp duplicate value tie breaker	15	service unavailable
tunnel rx recovery sccrq avp malformed bad length	15	service unavailable
tunnel rx recovery sccrq avp malformed truncated	15	service unavailable
tunnel rx recovery sccrq avp mismatched host name	15	service unavailable
tunnel rx recovery sccrq avp mismatched vendor name	15	service unavailable
tunnel rx recovery sccrq avp missing mandatory assigned tunnel id	15	service unavailable

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 24: Default L2TP Mappings *(continued)*

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
tunnel warmstart not operational	15	service unavailable
tunnel warmstart recovery error	15	service unavailable

Related Documentation

- Mapping Application Terminate Reasons and RADIUS Terminate Codes on page 121
- Configuring Custom Terminate Reason Mappings on page 148

PPP Terminate Reasons

Table 25 on page 141 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 © 2013, Juniper Networks, Inc.

Table 25: Default PPP Mappings

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
bundle fail link selection mismatch	10	NAS request
bundle fail local mped not set yet	10	NAS request

Copyright © 2013, Juniper Networks, Inc.

Table 25: Default PPP Mappings *(continued)*

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 25: Default PPP Mappings *(continued)*

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 25: Default PPP Mappings *(continued)*

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 25: Default PPP Mappings *(continued)*

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 25: Default PPP Mappings *(continued)*

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 25: Default PPP Mappings *(continued)*

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
osi admin disable	10	NAS request
osi link down	10	NAS request
osi max configure exceeded	10	NAS request
osi no local align npdu	10	NAS request
osi no peer align npdu	10	NAS request
osi no service	10	NAS request
osi peer renegotiate rx conf ack	10	NAS request
osi peer renegotiate rx conf nak	10	NAS request
osi peer renegotiate rx conf rej	10	NAS request
osi peer renegotiate rx conf req	10	NAS request
osi peer terminate code rej	10	NAS request
osi peer terminate term ack	10	NAS request
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

Copyright © 2013, Juniper Networks, Inc.

Table 25: Default PPP Mappings *(continued)*

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
subscriber mgr link interface not found	9	NAS error
subscriber mgr set state active failed	9	NAS error

Related Documentation

- Mapping Application Terminate Reasons and RADIUS Terminate Codes on page 121
- Configuring Custom Terminate Reason Mappings on page 148

Configuring Custom Terminate Reason Mappings

Junos OS supports default configuration mapping of terminate reasons for various protocols (AAA, DHCP, L2TP, and PPP) to RADIUS Acct-Terminate-Cause attributes. When a AAA, DHCP, L2TP, or PPP session is terminated, the router logs a message for the internal terminate reason and logs another message for the RADIUS Acct-Terminate-Cause attribute (RADIUS attribute 49). RADIUS attribute 49 is also included in RADIUS Acct-Off and Acct-Stop messages.

You can create customized mappings between a terminate reason and a RADIUS Acct-Terminate-Cause attribute to provide different information about the cause of a termination.

To configure customized mappings between a terminate reason and a RADIUS Acct-Terminate-Cause attribute:

1. Edit the **access** hierarchy.

   [edit]
   user@host# **edit access**

2. Edit the **terminate-code** statement.

   NOTE: **Terminate 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 | shutdown) | dchp | l2tp | ppp) that you want to modify.

   [edit access terminate-code]
   user@host# **set** ***protocol-option***

4. Specify the terminate reason that you want to modify.

   [edit access terminate-code *protocol-option*]
   user@host# **set** ***term-reason***

Copyright © 2013, Juniper Networks, Inc.

NOTE: Attempts to set a terminate reason mapping to its default value are rejected by the CLI.

5. Specify the RADIUS termination cause value (from 1 through 4294967295) that you want to use for the termination reason.

   [edit access terminate-code *protocol-option term-reason*]
   user@host# **set radius** ***term-cause***

NOTE: Deleting a customized mapping restores the default.

Related Documentation

- Mapping Application Terminate Reasons and RADIUS Terminate Codes on page 121
- AAA Terminate Reasons on page 123
- DHCP Terminate Reasons on page 124
- L2TP Terminate Reasons on page 124
- PPP Terminate Reasons on page 140

Configuring an Access Profile for Subscriber Management

Access profiles enable you to specify subscriber access authentication and accounting parameters. After access profiles are created, you can attach them at the **[edit logical-systems** ***logical-system-name*** **routing-instances** ***routing-instance-name*****]** hierarchy level or for use in automatically configuring VLANs or stacked VLANs at the **[edit interfaces** ***interface-name*** **auto-configure vlan-ranges]** or **[edit interfaces** ***interface-name*** **auto-configure stacked-vlan-ranges]** hierarchy levels.

To configure an access profile:

1. Edit the access stanza.

   [edit]
   user@host# **edit access**

2. Specify an existing or new access profile name.

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

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

Related Documentation

- Attaching Access Profiles on page 150
- Configuring Dynamic Authentication for VLAN Interfaces on page 728
- **profile on page 1899**

Attaching Access Profiles

After you have created the access profile that specifies the subscriber access management authentication and accounting parameters, you can attach the profile. Subscriber access management supports attaching access profiles at the following hierarchy levels:

- **[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.

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

Related Documentation

- AAA Service Framework Overview on page 21

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—192.0.43.10 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 at different levels of granularity: 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 © 2013, 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.

For IPv4 name server addresses, you can use either of two statements to configure the address. Addresses configured with the **domain-name-server-inet** statement take precedence over addresses configured with the **domain-name-server** statement.

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 151
- DHCP Attributes for Address-Assignment Pools on page 173
- Configuring DHCP Client-Specific Attributes on page 172

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 150.

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.

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 172.16.25.31
   user@host# set domain-name-server-inet 172.16.25.100
   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 172.20.10.01
   user@host# set domain-name-server-inet 172.20.100.100
   user@host# set domain-name-server-inet6 2001:db8:85a3::8a2e:370:ac81
   user@host# set domain-name-server-inet6 2001:db8:85a3::8a2e:370:71bfd
   ```

Related Documentation

- DNS Name Server Address Overview on page 150
- DHCP Attributes for Address-Assignment Pools on page 173
- Configuring DHCP Client-Specific Attributes on page 172

Copyright © 2013, Juniper Networks, Inc.

Understanding Session Options for Subscriber Access

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. 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 and downstream to the user. Control traffic is ignored. The subscriber is not considered idle as long as data traffic is detected in either direction.

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.

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.

Copyright © 2013, Juniper Networks, Inc.

- 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).

When using the idle timeout for dynamic VLAN removal, keep the following in mind:

- The idle timeout period begins after a dynamic subscriber VLAN interface is created or traffic activity stops on a dynamic subscriber VLAN interface.
- If a new client session is created or a client session is reactivated successfully, the client idle timeout resets.
- The removal of inactive subscriber VLANs functions only with VLANs that have been authenticated.

Related Documentation

- RADIUS IETF Attributes Supported by the AAA Service Framework on page 91
- Configuring Subscriber Session Options on page 154
- Removing Inactive Dynamic Subscriber VLANs on page 155

Configuring Subscriber Session Options

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. The subscriber session timeouts apply to both L2TP-tunneled and PPP-terminated subscriber sessions.

NOTE: **To configure the timeout attributes in RADIUS, refer to the documentation for your RADIUS server.**

To configure limitations on subscriber sessions:

1. Edit session options for the router access profile.

   [edit]
   user@host# **edit access profile** ***profile-name*** **session-options**

Copyright © 2013, Juniper Networks, Inc.

2. Configure the maximum period a subscriber session can be active.

   [edit access profile *profile-name* session-options]
   user@host# **set client-session-timeout** ***minutes***

3. Configure the maximum period a subscriber session can be idle.

   [edit access profile *profile-name* session-options]
   user@host# **set client-idle-timeout** ***minutes***

For example, to configure a client session timeout of 2 hours and an idle timeout of 15 minutes in the **acc-prof** profile:

```
[edit]
access {
  profile {
    acc-prof {
      session-options {
        client-session-timeout 120;
        client-idle-timeout 15;
      }
    }
  }
}
```

Related Documentation

- Understanding Session Options for Subscriber Access on page 153
- client-idle-timeout on page 1511
- client-session-timeout on page 1512

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: **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 153
- Configuring Subscriber Session Options on page 154
- **client-idle-timeout on page 1511**

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 156

Testing a Subscriber AAA Configuration

Purpose Display the AAA attributes that subscriber management assigns to the subscriber during login.

The following example tests the AAA configuration for a PPP subscriber. You can use the ***test aaa dhcp user*** command to perform a similar test for DHCP subscribers and the ***test aaa authd-lite user*** command to test authd-lite subscribers.

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 © 2013, Juniper Networks, Inc.

Action `user@host>` **test aaa ppp user thomastank password 00N15&**

```
Authentication Grant
************User Attributes***********
     User Name -                            thomastank
     Client IP Address -                    192.168.1.1
     Client IP Netmask -                    255.255.0.0
     Virtual Router Name -                  default
     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
     Framed Pool -                          addr_pool2
     Session Timeout -                      0
     Idle Timeout -                         0
     Service Type -                         0
     Client Ipv6 Address -                  ::
     Client Ipv6 Mask -                     null
     Framed Ipv6 Prefix -                   ::/0
     Framed Ipv6 Pool -                     not-set
     Nas Ipv6 Address -                     ::
     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 -                        00N15&
     CHAP Password -                        NULL
     NAS Ip Address -                       0.0.0.0
     Agent Remote Id -                      not-set
     NAS Port -                             0
     NAS Port Type -                        5
Client Session Activate request sent
Client Session Activated
     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
     LI Action -                            0
     LI Intercpet Id -                      0
     Med Ipaddress -                        0.0.0.0
     Med Port Number -                      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
     IGMP Version -                         IGMP Version not set
     MLD Version -                          MLD Version not set
     IGMP Immediate Leave -                 disabled
     MLD Immediate Leave -                  disabled
     IPv6 Ingress Policy Name -             not set
     IPv6 Egress Policy Name -              not set
     Cos Parameter Type -                   not-set
     Cos  Scheduler Parameter Type -        not-set
```

```
     Acct Session ID-                        8
     Acct Interim Interval -                 0
     Acct Type -                             0
     Ingress Statistics                      disabled
     Egress Statistics                       disabled
****Pausing 10 seconds before disconnecting the test user*********
Logging out subscriber
Test complete. Exiting
```

If you specify the DSL Forum Agent-Remote-Id (VSA 26-2), the output includes the specified value:

```
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
     ...
```

If you do not specify the VSA, then the Agent-Remote-Id value is shown as **not-set**.

Related Documentation

- AAA Configuration Testing and Troubleshooting on page 156

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 *Junos OS System Log Messages Reference*.)

Copyright © 2013, 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.

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 159.

2. Configure the number and size of trace logs.

   See "Configuring the Number and Size of Subscriber Management Database Log Files" on page 159.

3. Configure user access to trace logs.

   See "Configuring Access to the Subscriber Management Database Log File" on page 160.

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 160.

5. Configure flags to specify which events are logged.

   See "Configuring the Subscriber Management Database Tracing Flags" on page 161.

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 158

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 158

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 158

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.

 Copyright © 2013, Juniper Networks, Inc.

```
[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 158

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
```

Related Documentation

- Tracing Subscriber Management Database Operations for Subscriber Access on page 158

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 *Junos OS System Log Messages Reference*.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure 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 164.

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 162.

3. Configure user access to trace logs.

   See "Configuring Access to the Subscriber Management Session Database Replication Log File" on page 163.

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 163.

5. Configure flags to specify which events are logged.

   See "Configuring the Subscriber Management Session Database Replication Tracing Flags" on page 164.

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 © 2013, Juniper Networks, Inc.

Related Documentation

- Tracing Subscriber Management Session Database Replication Operations for Subscriber Access on page 161

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 161

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 161

Copyright © 2013, Juniper Networks, Inc.

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 161

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.

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 161

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:

 Copyright © 2013, Juniper Networks, Inc.

user@host> **clear network-access aaa subscriber**

- To clear AAA accounting statistics:

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

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

  user@host> **clear network-access aaa statistics address-assignment client**

- To clear AAA address-assignment pool statistics:

  user@host> **clear network-access aaa statistics address-assignment pool** ***pool-name***

- To clear AAA authentication statistics:

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

Related Documentation

- *Junos OS Operational Mode Commands*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 3

Configuring Address-Assignment Pools for Subscriber Access

- Address-Assignment Pools Overview on page 167
- Configuring Address-Assignment Pools on page 168
- Configuring an Address-Assignment Pool Name and Addresses on page 169
- Configuring a Named Address Range for Dynamic Address Assignment on page 169
- Configuring Address-Assignment Pool Linking on page 170
- Configuring Static Address Assignment on page 171
- Configuring an Address-Assignment Pool for Router Advertisement on page 171
- Configuring Address-Assignment Pool Usage Threshold Traps on page 172
- Configuring DHCP Client-Specific Attributes on page 172
- DHCP Attributes for Address-Assignment Pools on page 173
- Address-Assignment Pools Licensing Requirements on page 175
- Tracing General Authentication Service Processes on page 175

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 © 2013, 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 the primary pool is fully allocated, the router or switch automatically switches to the linked, or secondary, pool and begins allocating addresses from that pool.

You can also explicitly identify that an address-assignment pool is used for ND/RA.

Related Documentation

- Configuring Address-Assignment Pools on page 168
- DNS Address Assignment Precedence on page 44
- Address-Assignment Pools Licensing Requirements on page 175
- Example: Configuring an Address-Assignment Pool on page 195
- *Configuring an Extended DHCP Server with DHCPv6 on EX Series Switches (CLI Procedure)*

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: **You cannot use address-assignment pools with the J Series Services Routers DHCP server. Also, 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 169.

2. (Optional) Configure named ranges (subsets) of addresses.

   See "Configuring a Named Address Range for Dynamic Address Assignment" on page 169.

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 170.

Copyright © 2013, Juniper Networks, Inc.

4. (Optional) Create static address bindings (IPv4 only).

   See "Configuring Static Address Assignment" on page 171.

5. (Optional) Configure attributes for DHCP clients.

   See "Configuring DHCP Client-Specific Attributes" on page 172.

6. (Optional) Specify that the address-assignment pool is used for router advertisement.

   See "Configuring an Address-Assignment Pool for Router Advertisement" on page 171.

Related Documentation

- Address-Assignment Pools Overview on page 167
- Address-Assignment Pools Licensing Requirements on page 175
- Example: Configuring an Address-Assignment Pool on page 195

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 2008:2009::/32
   ```

Related Documentation

- Address-Assignment Pools Overview on page 167
- Configuring Address-Assignment Pools on page 168

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 2008:2010:2011:0100::/64 high
     2008:2010:2011:ffff::/64
   user@host# set range fiber-east prefix-length 48
   ```

Related Documentation

- Address-Assignment Pools Overview on page 167
- Configuring Address-Assignment Pools on page 168

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 is 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. There is no limit to the number of linked pools in a chain. However, you cannot create multiple links to or from the same pool—a pool can be linked to only one secondary pool, and a secondary pool can be linked from only one primary pool. Also, two linked primary and secondary pools must be of the same family type, either IPv4 or IPv6.

To link an address-assignment pool to a secondary pool:

1. Specify the name of the primary address-assignment pool.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit access]
user@host# edit address-assignment pool pool-name
```

2. Configure the secondary pool to which the primary pool will be linked.

```
[edit access address-assignment pool isp_1]
user@host# set link pool-name
```

Related Documentation

- Address-Assignment Pools Overview on page 167
- Address-Assignment Pools Licensing Requirements on page 175
- Example: Configuring an Address-Assignment Pool on page 195

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 90:00:00:01:00:01 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 90:00:00:01:00:01
  ip-address 192.168.44.12
```

Related Documentation

- Address-Assignment Pools Overview on page 167
- Configuring Address-Assignment Pools on page 168

Configuring an Address-Assignment Pool for Router Advertisement

You can create an address-assignment pool that is explicitly used for router advertisement address assignment. You populate the address-assignment pool using the standard procedure, but you additionally specify that the pool is used for router advertisement.

To create an address-assignment pool that is used for router advertisement:

1. Configure the IPv6 address-assignment pool.

   See "Configuring Address-Assignment Pools" on page 168

2. Specify that the address-assignment pool is used for router advertisement.

```
[edit access address-assignment]
user@host# set neighbor-discovery-router-advertisement chi-fiber-ra
```

Related Documentation

- Address-Assignment Pools Overview on page 167
- Example: Configuring an Address-Assignment Pool on page 195

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 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 167
- Example: Configuring an Address-Assignment Pool on page 195

Configuring DHCP Client-Specific Attributes

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

 Copyright © 2013, Juniper Networks, Inc.

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 173 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 isp_1 family inet**

2. Configure optional DHCP client attributes.

   [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 150.

Related Documentation

- Address-Assignment Pools Overview on page 167
- Configuring Address-Assignment Pools on page 168
- DHCP Attributes for Address-Assignment Pools on page 173

DHCP Attributes for Address-Assignment Pools

Table 26 on page 173 describes the DHCP client attributes that you can use with the **dhcp-attributes** statement when you configure address-assignment pools. Table 27 on page 174 describes the DHCPv6 client attributes for configuring IPv6 address-assignment pools.

Table 26: 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.	–

Table 26: DHCP Attributes *(continued)*

Attribute	Description	DHCP Option
maximum-lease-time	Maximum lease time allowed by the DHCP server.	51
name-server	IP address of DNS server to which clients can send DNS queries.	6
netbios-node-type	NetBIOS node type.	46
option	User-defined options.	–
option-match	Option 82 value is mapped to named address range.	–
router	IP address for routers on the subnetwork.	3
server-identifier	IP address used as the DHCP source address	54
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 27: DHCPv6 Attributes

Attribute	Description	DHCPv6 Option
dns-server	IPv6 address of DNS server to which clients can send DNS queries.	23
grace-period	Grace period offered with the lease.	–
maximum-lease-time	Maximum lease time allowed by the DHCP server.	–
option	User-defined options.	–
preferred-lifetime	Length of time, in seconds, that the DHCP server keeps the IPv6 prefix active. The DHCPv6 server sends this attribute to the client (router).	–
sip-server-address	IPv6 address of SIP outbound proxy server.	22
sip-server-domain-name	Domain name of the SIP outbound proxy server.	21

Copyright © 2013, Juniper Networks, Inc.

Table 27: DHCPv6 Attributes *(continued)*

Attribute	Description	DHCPv6 Option
t1-percentage	Timer that, when it expires, the client (router) sends a request to the originating DHCP server requesting its lease (IP configuration) to be extended. If the lease period is limitless, the client's lease never expires and no request is required. However, because most leases are time-limited, the router monitors this timer and sends a request to the originating DHCP server to ensure that its lease is extended and normal operation continues. T1-percentage is specified as a percentage of the elapsed time for the preferred-lifetime value. This value is sent from the DHCP server to the router.	–
t2-percentage	Timer that, when it expires, the client (router) sends a request to any available DHCP server requesting its lease (IP configuration) to be extended. T2-percentage is sent from the DHCP server to the router. You specify t2-percentage as a percentage of the value of preferred-lifetime.	–
valid-lifetime	Length of time, in seconds, that the DHCP server keeps the IPv6 prefix valid.	–

Related Documentation

- Address-Assignment Pools Overview on page 167
- Configuring Address-Assignment Pools on page 168
- Configuring DHCP Client-Specific Attributes on page 172
- **dhcp-attributes (Address-Assignment Pools) on page 1542**

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*

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 *Junos OS System Log Messages Reference*.
- 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 176
- Configuring the Number and Size of General Authentication Service Processes Log Files on page 176
- Configuring Access to the Log File on page 177
- Configuring a Regular Expression for Lines to Be Logged on page 177
- Configuring the Trace Operation on page 178

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

 Copyright © 2013, Juniper Networks, Inc.

existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output, by including the **files** and **size** options with the **traceoptions** statement.

  [edit system processes general-authentication-service traceoptions]
  user@host# **set file aap_logfile_1 files 20 size 2097152**

Configuring Access to the 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 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**

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 4

Configuring Domain Maps for Subscriber Access

- Domain Mapping Overview on page 180
- Configuring a Domain Map on page 181
- Specifying an Access Profile in a Domain Map on page 182
- Specifying a Dynamic Profile in a Domain Map on page 183
- Specifying an Address Pool in a Domain Map on page 184
- Specifying an AAA Logical System/Routing Instance in a Domain Map on page 184
- Specifying a Target Logical System/Routing Instance in a Domain Map on page 186
- Configuring Domain Name Usage for Domain Maps on page 187
- Specifying Domain Name Delimiters on page 187
- Specifying the Parsing Direction for Domain Names on page 188
- Enabling Domain Name Stripping on page 189
- Specifying a Tunnel Profile in a Domain Map on page 189
- Specifying a Tunnel Switch Profile in a Domain Map on page 190
- Configuring PADN Parameters for a Domain Map on page 190
- Verifying and Managing Domain Map Configuration on page 191

Copyright © 2013, 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 **xyz.com**. The options and parameters in that domain map are then applied when subscribers with the specified domain name (for example, **bob@xyz.com**, **raj@xyz.com**, and **juan@xyz.com**) request an AAA service.

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 and are not supported in current Junos OS releases.**

Table 28 on page 180 describes the access options and parameters you can configure in the domain map.

Table 28: 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 name rules	Rules for domain 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.

Copyright © 2013, Juniper Networks, Inc.

Table 28: Domain Map Options and Parameters *(continued)*

Option	Description
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.

Default Domain Map

You can configure a default domain map that the router uses for subscribers whose domain name does not explicitly match any existing domain map. The router also uses the default domain map when a subscriber username does not include a domain name.

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 match a subscriber request to a domain map, the router then uses the rules specified in the default domain map configuration to handle the subscriber request.

Related Documentation

- Configuring a Domain Map on page 181

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 **xyz.com**:

     [edit access]
     user@host# **edit domain map xyz.com**

   - To create a default domain map to be mapped to subscribers with non-matching domain names and subscribers without domain names:

     [edit access]
     user@host# **edit domain map default**

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 182.

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 183.

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 184.

5. (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 Name Usage for Domain Maps" on page 187.

6. (Optional) Assign a tunnel profile that provides tunnel definitions for the domain map.

   See "Specifying a Tunnel Profile in a Domain Map" on page 189.

7. (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 190.

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

   See "Configuring PADN Parameters for a Domain Map" on page 190.

Related Documentation

- Domain Mapping Overview on page 180
- Verifying and Managing Domain Map Configuration on page 191

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 29 on page 182.

Table 29: 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

Copyright © 2013, Juniper Networks, Inc.

Table 29: Precedence Rules for Applying Access Profiles *(continued)*

Precedence (High to Low)	How the Access Profile Is Applied
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 180
- Configuring a Domain Map on page 181

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 30 on page 183.

Table 30: 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 180
- Configuring a Domain Map on page 181

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 31 on page 184.

Table 31: 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 180
- Configuring a Domain Map on page 181

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

Copyright © 2013, Juniper Networks, Inc.

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 **xyz.com** are now authenticated by a RADIUS server in a particular non-default AAA context.

NOTE: Changing the AAA context does not change the subscriber context. You use the target-logical-system command to explicitly configure the logical system/routing instance for subscribers.

To configure the default logical system and a non-default routing instance for AAA requests:

1. Specify the domain map you want to configure.

   [edit access]
   user@host# **edit domain map** ***domain-map-name***

2. Specify the non-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***

NOTE: Subscriber management is supported in the default logical system only. The following procedure, which describes configuring a non-default logical system, is for future extensions of subscriber management and is not supported in current Junos OS releases.

To configure a non-default logical system in which you want the **authd** daemon to send AAA requests:

1. Specify the domain map you want to configure.

   [edit access]
   user@host# **edit domain map** ***domain-map-name***

2. Specify the logical system and optionally the non-default routing instance for AAA requests.

   - To configure a non-default logical system and default routing instance for AAA requests:

     [edit access domain map *domain-map-name*]
     user@host# **set** aaa-logical-system ***logical-system-name***

   - To configure a non-default logical system and a non-default routing instance for AAA requests:

     [edit access domain map *domain-map-name*]
     user@host# **set** aaa-logical-system ***logical-system-name*** aaa-routing-instance ***routing-instance-name***

Related Documentation

- Domain Mapping Overview on page 180
- Configuring a Domain Map on page 181
- Specifying a Target Logical System/Routing Instance in a Domain Map on page 186

Specifying a Target Logical System/Routing Instance in a Domain Map

By default, the router places a subscriber in the logical system/routing instance of the interface on which the subscriber negotiations start. Subscriber management can then use the authentication server or a domain map to change the subscriber's logical system/routing instance.

To use the domain map method, you configure the domain map to specify the target logical system and routing instance for the subscriber's interface. You can optionally configure the domain map to use the default logical system and a specific non-default routing instance.

To configure a default target logical system and a non-default routing instance 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 non-default target routing instance. The target logical system is automatically set to the default.

   [edit access domain map *domain-map-name*]
   user@host# **set** target-routing-instance ***routing-instance-name***

NOTE: Subscriber management is supported in the default logical system only. The following procedure, which describes configuring a non-default target logical system, is for future extensions of subscriber management and is not supported in current Junos OS releases.

To configure a non-default target logical system 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 logical system and, optionally, the non-default target routing instance for the subscriber's interface.

   - To configure a non-default target logical system and default target routing instance:

     [edit access domain map *domain-map-name*]
     user@host# **set** target-logical-system ***logical-system-name***

   - To configure a non-default target logical system and a non-default target routing instance:

Copyright © 2013, Juniper Networks, Inc.

```
[edit access domain map domain-map-name]
user@host# set target-logical-system logical-system-name target-routing-instance
  routing-instance-name
```

Related Documentation

- Domain Mapping Overview on page 180
- Configuring a Domain Map on page 181

Configuring Domain Name Usage for Domain Maps

You can configure how the router determines domain names 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 name delimiters and to specify the direction the router uses to parse domain names. At the domain map level, you can enable domain name stripping. Domain name stripping specifies that the router remove the domain 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) Configure the domain name delimiters you want the router to recognize for domain maps.

   See "Specifying Domain Name Delimiters" on page 187.

2. (Optional) Configure the parse direction you want the router to use when determining domain names for domain maps.

   See "Specifying the Parsing Direction for Domain Names" on page 188.

3. (Optional) Configure the router to remove the domain name from usernames in the domain map before using AAA services.

   See "Enabling Domain Name Stripping" on page 189.

Related Documentation

- Domain Mapping Overview on page 180
- Configuring a Domain Map on page 181

Specifying Domain Name Delimiters

A delimiter is the character that separates a subscriber username from the domain name. Delimiters are commonly used for domain name parsing or stripping. You can specify a maximum of eight delimiters that the router uses to recognize domain names for a domain map. If you do not configure any delimiters, the router uses the @ character by default.

For example, your network might include the subscribers **bob@abc.com**, **pete!xyz.com**, and **maria\pqr.com**. In this case, you configure the router to recognize the characters **@**, **!**, and **** as delimiters.

Keep the following guidelines in mind when specifying delimiters:

- You cannot use the semicolon (;) as a delimiter.
- If you configure optional 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 delimiters and then unconfigure them, the router sets the domain map delimiter back to the default @ character.

To configure domain 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 delimiters. Do not include spaces between the delimiters.

   [edit access domain]
   user@host# **set delimiter** [*delimiter-character*]

Related Documentation

- Configuring Domain Name Usage for Domain Maps on page 187

Specifying the Parsing Direction for Domain Names

You can specify the direction in which the router performs the parsing operation it uses to identify subscriber domain 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 parsing direction you use is important when there are nested domain names. For example, for the username **user1@abc.com@xyz.com**, right-to-left parsing produces a domain name of **xyz.com**. For the same username, left-to-right parsing produces a domain name of **abc.com@xyz.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.

   [edit access domain]
   user@host# **set parse-direction** (left-to-right | right-to-left)

Related Documentation

- Configuring Domain Name Usage for Domain Maps on page 187

Copyright © 2013, Juniper Networks, Inc.

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@xyz.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 Name Usage for Domain Maps on page 187

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 180
- Configuring a Domain Map on page 181
- Configuring a Tunnel Profile for Subscriber Access on page 403

Copyright © 2013, Juniper Networks, Inc.

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-Group attribute (VSA 26-64) takes precedence over the tunnel switch profile specified in the domain map.

To include a tunnel switch profile in a domain map:

1. Specify the domain map you want to configure.

   [edit access]
   user@host# edit domain map *domain-map-name*

2. Specify the tunnel switch profile you want to include in the domain map.

   [edit access domain map *domain-map-name*]
   user@host# set tunnel-switch-profile *profile-name*

Related Documentation

- Domain Mapping Overview on page 180
- Configuring a Domain Map on page 181
- Configuring L2TP Tunnel Switching on page 432

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*

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Domain Mapping Overview on page 180
- Configuring a Domain Map on page 181

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 180
- Configuring a Domain Map on page 181

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 5

AAA and Remote Subscriber Access Configuration Examples

- Example: Configuring RADIUS-Based Subscriber Authentication and Accounting on page 193
- Example: Configuring an Address-Assignment Pool on page 195

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;
      retry 3;
      secret &tIUEI*7688+;
      source-address 192.168.1.100;
      timeout 45;
   }
   192.168.1.251 {
      port 1812;
      accounting-port 1813;
      retry 3;
      secret $Dyu*UY(877-;
      source-address 192.168.1.100;
      timeout 30;
   }
   192.168.1.252 {
      port 1812;
      secret $Dyu*UY(877-;
   }
}
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];
                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 {
```

 Copyright © 2013, Juniper Networks, Inc.

```
            unnumbered-address lo0.0;
          }
        }
      }
    }
```

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers on page 23

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.net {
        hardware-address 90:00:00:01:00:01;
        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 2008:2009:2010::/48;
      range fiber3 {
        low 2008:2009:2010::1/64;
        high 2008:2009:2010::5/64;
```

```
        }
      }
    }
  }
```

This example creates an IPv4 address-assignment pool named **isp-1**, which contains two named address ranges, **southeast** and **northeast**. The address-assignment pool also contains a static binding for client **host sval6.boston.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 167
- Configuring Address-Assignment Pools on page 168
- Configuring an Address-Assignment Pool for Router Advertisement on page 171

Copyright © 2013, Juniper Networks, Inc.

PART 3

DHCP Local Server for Subscriber Access

- DHCP Local Server Overview on page 199
- Configuring DHCP Local Server on page 211
- DHCP Local Server Examples on page 263

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 6

DHCP Local Server Overview

- Extended DHCP Local Server Overview on page 200
- DHCPv6 Local Server Overview on page 204
- DHCP Local Server Handling of Client Information Request Messages on page 205
- Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview on page 207
- Port Number Requirements for DHCP Firewall Filters on page 208

Extended DHCP Local Server Overview

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

The extended DHCP local server enhances traditional DHCP server operation by utilizing centralized address-assignment pools. The address-assignment pools are managed by the **authd** process, independently of the DHCP local server, and can be shared by different client applications.

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. See "DHCPv6 Local Server Overview" on page 204 for information about the DHCPv6 local server feature.

NOTE: **You cannot configure the extended DHCP local server and extended DHCP relay on the same interface.**

To configure the extended DHCP local server on the router (or switch), you include the **dhcp-local-server** statement at the **[edit system services]** hierarchy level. See the "[edit system services dhcp-local-server] Hierarchy Level" on page 1407 for the complete DHCP local server syntax.

This overview covers:

- Interaction Among the DHCP Client, Extended DHCP Local Server, and Address-Assignment Pools on page 200
- Providing DHCP Client Configuration Information on page 201
- Minimal Configuration for Clients on page 202
- DHCP Local Server and Address-Assignment Pools on page 202
- DHCP Liveness Detection on page 203

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 the software installation is on a router or a switch. Technically, the codes operates in the same manner, regardless of the hardware platform. However, there are some difference in the details of usage.

 Copyright © 2013, Juniper Networks, Inc.

- On routers—In a typical carrier edge network configuration, the DHCP client is on the subscriber's computer, 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.

IThe 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.

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, 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 must configure the local address-assignment pool to provide the configuration 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 32 on page 202 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 90 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 32: 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)
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

 Copyright © 2013, Juniper Networks, Inc.

address and configuration information reside in centralized address-assignment pools. The centralized 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. 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.**

DHCP Liveness Detection

Liveness detection for DHCP subscriber IP (or DHCP client IP) sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients 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. You can configure

NOTE: **DHCP liveness detection either globally or per DHCP group.**

Related Documentation

- Address-Assignment Pools Overview on page 167
- Configuring Address-Assignment Pools on page 168
- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 215
- Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview on page 207
- Using External AAA Authentication Services with DHCP on page 214
- Use of DHCP Option 50 and DHCPv6 IA_NA Option to Request a Specific IP Address on page 230
- Graceful Routing Engine Switchover on page 254

Copyright © 2013, Juniper Networks, Inc.

- Subscriber Management Unified ISSU Support on page 7
- Tracing Extended DHCP Operations on page 255
- Verifying and Managing DHCP Local Server Configuration on page 253
- Example: Minimum Extended DHCP Local Server Configuration on page 263
- Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 263
- Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine on page 265

DHCPv6 Local Server Overview

The DHCPv6 local server enhances the extended DHCP local server by providing support for IPv6. 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.

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

NOTE: **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.**

Based on the attributes that the DHCPv6 local server provides, RADIUS returns the information listed in Table 33 on page 204 to configure the client:

Table 33: 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

The DHCPv6 local server is compatible with the extended DHCP local server and the extended DHCP relay agent, and can be enabled on the same interface as either the extended DHCP local server or DHCP relay agent.

Copyright © 2013, Juniper Networks, Inc.

The DHCPv6 local server provides many of the same features as the extended 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

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. See the "[edit system services dhcp-local-server] Hierarchy Level" on page 1407 for the complete DHCP local server syntax, including the DHCPv6 syntax.

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 200
- Using External AAA Authentication Services with DHCP on page 214
- Grouping Interfaces with Common DHCP Configurations on page 217
- Group-Specific DHCP Local Server Options on page 219
- Overriding Default DHCP Local Server Configuration Settings on page 220
- Configuring Passwords for Usernames on page 238
- Creating Unique Usernames for DHCP Clients on page 239
- Use of DHCP Option 50 and DHCPv6 IA_NA Option to Request a Specific IP Address on page 230
- Verifying and Managing DHCPv6 Local Server Configuration on page 253
- Example: Extended DHCPv6 Local Server Configuration on page 264

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. Include the **process-inform**

statement at the **[edit system services dhcp-local-server overrides]** or **[edit system services dhcp-local-server dhcpv6 overrides]** hierarchy level.

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.

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.

NOTE: **PPP interfaces are not supported on EX Series switches.**

When DHCPv6 is configured over PPP interfaces, the PPP RADIUS authentication data can be used to select the pool from which the response information is taken. Additionally other RADIUS attributes can also be inserted into the DHCPv6 reply message. If an overlap exists between RADIUS attributes and local pool attributes, the RADIUS values are used instead of the local configuration data. If no RADIUS information is received from the underlying PPP interface, then the behavior is the same as described previously for non-PPP interfaces.

Related Documentation

- Overriding Default DHCP Local Server Configuration Settings on page 220
- Enabling Processing of Client Information Requests on page 227

Copyright © 2013, Juniper Networks, Inc.

Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview

The router's DHCP support enables you to attach a dynamic profile to a DHCP subscriber interface. When a DHCP subscriber logs in, the router instantiates the specified dynamic profile and then applies the services defined in the profile to the interface.

You can attach dynamic profiles to all interfaces or you can specify a particular group of interfaces to which the profile is attached. Both the DHCP local server and the DHCP relay agent support the attachment of dynamic profiles to interfaces.

You can enable the following optional features when the dynamic profile is attached. The two options cannot be used together.

- Enable multiple DHCP subscribers to share the same VLAN logical interface. The firewall filters, CoS schedulers, and IGMP configuration of the clients are merged.
- Specify the primary dynamic profile that is instantiated when the first subscriber logs in.

Multiple DHCP Subscribers Sharing the Same VLAN Logical Interface

The **aggregate-clients** statement specifies that the router merge the firewall filters, CoS schedulers, and IGMP configuration of multiple DHCP clients that are on the same VLAN logical interface (for example, multiple clients belonging to the same household). You can configure the aggregate-clients support for all interfaces or for a group of interfaces. The **aggregate-clients** statement provides the option of either merging (chaining) or replacing software components for each client.

By default, the feature is disabled and a single DHCP client is allowed per VLAN when a dynamic profile is associated with the VLAN logical interface.

When you specify the **merge** option, the router aggregates the software components for multiple subscribers as follows:

- Firewall filters—The filters are chained together using the precedence as the order of execution. If the same firewall filter is attached multiple times, the filter is executed only once.
- CoS schedulers—The different CoS schedulers are merged as if the scheduler map has multiple schedulers. The merge operation for the individual traffic-control-profiles parameters (shaping-rate, delay-buffer-rate, guaranteed-rate) preserves the maximum value for each parameter.
- IGMP configuration—The current IGMP configuration is replaced with the configuration of the newest DHCP client.

When you specify the **replace** option, the entire logical interface is replaced whenever a new client logs in to the network using the same VLAN logical interface. For example, if a customer subscribes to voice, video, and data services on the network, when a voice client logs in, instead of applying a specific voice filter for only that service, the entire voice, video, and data filter chain is applied.

Copyright © 2013, Juniper Networks, Inc.

NOTE: You cannot use a dynamic demux interface to represent multiple subscribers in a dynamic profile attached to an interface. One dynamic demux interface represents one subscriber. Do not configure the **aggregate-clients** option when attaching a dynamic profile to a demux interface for DHCP.

Primary Dynamic Profile

The **use-primary** option enables you to specify the primary dynamic profile that 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.

This feature can conserve logical interfaces in a network where dynamic IP demux interfaces are used to represent subscribers. To conserve interfaces, make sure the primary profile that you specify does not create a demux interface, but provides the initial policies for the primary interface subscriber.

Related Documentation

- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

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, 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.

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 © 2013, Juniper Networks, Inc.

For complete information about configuring firewall filters in general, see *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*.

Related Documentation

- Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine on page 265
- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276
- Dynamic Firewall Filters Overview on page 1132

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 7

Configuring DHCP Local Server

- DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 212
- Guidelines for Configuring Support for DHCP Duplicate Clients on page 213
- Configuring DHCP Duplicate Client Support on page 214
- Using External AAA Authentication Services with DHCP on page 214
- Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 215
- Grouping Interfaces with Common DHCP Configurations on page 217
- Guidelines for Configuring Interface Ranges on page 218
- Group-Specific DHCP Local Server Options on page 219
- Overriding Default DHCP Local Server Configuration Settings on page 220
- Specifying the Maximum Number of DHCP Clients Per Interface on page 221
- DHCP Access, Access-Internal, and Destination Route Suppression on page 222
- Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default on page 223
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 223
- DHCP Auto Logout Overview on page 225
- Automatically Logging Out DHCP Clients on page 226
- Enabling Processing of Client Information Requests on page 227
- Enabling DHCPv6 Rapid Commit Support on page 228
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229
- Use of DHCP Option 50 and DHCPv6 IA_NA Option to Request a Specific IP Address on page 230
- Multiple Address Assignment for DHCPv6 Clients on page 230
- Specifying the Delegated Address Pool for IPv6 Prefix Assignment on page 231
- Subscriber Binding Retention During Interface Delete Events on page 232
- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 232
- Verifying and Managing the DHCP Maintain Subscribers Feature on page 233

Copyright © 2013, Juniper Networks, Inc.

- Clearing DHCP Bindings for Subscriber Access on page 234
- Configuring the Router to Strictly Enforce the Subscriber Scaling License on page 235
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236
- Configuring Passwords for Usernames on page 238
- Creating Unique Usernames for DHCP Clients on page 239
- Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients on page 241
- Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 246
- Configuring Deletion of the Client When Dynamic Reconfiguration Fails on page 247
- Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 247
- Configuring a Token for DHCP Local Server Authentication on page 248
- Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings on page 248
- Preventing Binding of Clients That Do Not Support Reconfigure Messages on page 249
- DHCP Liveness Detection Overview on page 250
- Configuring Detection of DHCP Local Server Client Connectivity on page 251
- Verifying and Managing DHCP Local Server Configuration on page 253
- Verifying and Managing DHCPv6 Local Server Configuration on page 253
- Graceful Routing Engine Switchover on page 254
- Tracing Extended DHCP Operations on page 255

DHCP Duplicate Client Differentiation Using Client Subinterface Overview

In some network environments, client IDs and MAC addresses might not be unique, resulting in duplicate clients. For example, two network adapters might be manufactured with the same hardware address, resulting in a duplicate MAC address among the DHCP clients attached to the router (or switch). A duplicate DHCP client occurs when a client attempts to get a lease, and that client has the same client ID or the same MAC address as an existing DHCP client.

When DHCP server receives a request from a new client that has a duplicate ID or MAC address, DHCP server terminates the address lease for the existing client and returns the address to its original address pool. DHCP server then assigns a new address and lease to the new client.

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

 Copyright © 2013, Juniper Networks, Inc.

You can provide greater differentiation between duplicate clients by configuring DHCP to consider the client subinterface when duplicate clients occur. In this optional configuration, DHCP uniquely identifies:

- The subnet on which the client resides
- The subinterface on which the client resides
- The client within the subnet

Related Documentation

- Configuring DHCP Duplicate Client Support on page 214
- Guidelines for Configuring Support for DHCP Duplicate Clients on page 213

Guidelines for Configuring Support for DHCP Duplicate Clients

This topic describes the guidelines for configuring DHCP to include the client subinterface in order to distinguish between duplicate clients (clients with the same MAC address or client ID) in a subscriber access or in a DHCP-managed environment.

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

- The optional DHCP duplicate client support feature is used for DHCPv4 clients. For DHCPv6, client identification is independent of MAC address.
- For DHCP relay agent configuration:
  - DHCP relay must be configured to insert option 82, regardless of whether or not the incoming packet has option 82.
  - Option 82 must include the Agent Circuit ID suboption (suboption 1).
  - Option 82 must be the interface name, not the interface description.
  - DHCP server must echo option 82 in the server's reply. This is required because of the following:
    - The giaddr inserted by DHCP relay is the same for duplicate clients on different subinterfaces. The DHCP local server uses option 82 when allocating the IP address.
    - DHCP relay uses the echoed option 82 to learn the client subinterface and to construct the client key.
- For the Layer 3 wholesale model:
  - The wholesaler and retailer logical system/routing instances must have the same **duplicate-clients-on-interface** statement configuration.
  - For DHCP relay, the wholesaler and the retailer routing contexts must both be configured with the Agent Circuit ID suboption (suboption 1) in option 82.

Related Documentation

- DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 212
- Configuring DHCP Duplicate Client Support on page 214

Copyright © 2013, Juniper Networks, Inc.

Configuring DHCP Duplicate Client Support

You can optionally configure DHCP local server and DHCP relay to include a client subinterface when distinguishing between two clients that have the same MAC address or client ID. The configuration is a global setting for each logical system/routing instance.

To configure DHCP local server to include the client subinterface:

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

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

2. Configure the optional duplicate client support.

   [edit system services dhcp-local-server]
   user@host# **set** duplicate-clients-on-interface

To configure DHCP relay agent to include the client subinterface:

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

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

2. Configure the optional duplicate client support.

   [edit system services dhcp-relay]
   user@host# **set** duplicate-clients-on-interface

Related Documentation

- DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 212
- Guidelines for Configuring Support for DHCP Duplicate Clients on page 213

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

 Copyright © 2013, Juniper Networks, Inc.

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.
   - 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 238.

3. (Optional) Configure optional features to create a unique username.

   See "Creating Unique Usernames for DHCP Clients" on page 239.

Related Documentation

- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276
- DHCPv6 Local Server Overview on page 204
- DHCPv6 Relay Agent Overview on page 279

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]
     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**

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Address-Assignment Pools Overview on page 167
- Configuring Address-Assignment Pools on page 168
- Extended DHCP Local Server Overview on page 200
- Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 263

Grouping Interfaces with Common DHCP Configurations

You use the group feature to group together a set of interfaces and then apply a common DHCP configuration to the named interface group. The extended DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent all support interface groups.

The following steps create a DHCP local server group; the steps are similar for the DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent.

To configure a DHCP local server interface group:

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

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

2. Create the group and assign a name.

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

3. Specify the names of one or more interfaces on which the extended DHCP application is enabled. You can repeat the **interface** ***interface-name*** statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

   ```
   [edit system services dhcp-local-server group boston]
   user@host# set interface fe-1/0/1.1
   user@host# set interface fe-1/0/1.2
   ```

4. (Optional) You can use the **upto** option to specify a range of interfaces for a group.

   ```
   [edit system services dhcp-local-server group boston]
   user@host# set interface fe-1/0/1.3 upto fe-1/0/1.9
   ```

5. (Optional) You can use the **exclude** option to exclude a specific interface or a specified range of interfaces from the group. For example:

   ```
   [edit system services dhcp-local-server group boston]
   user@host# set interface fe-1/0/1.1 upto fe-1/0/1.102
   user@host# set interface fe-1/0/1.6 exclude
   user@host# set interface fe-1/0/1.70 upto fe-1/0/1.80 exclude
   ```

Related Documentation

- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276

- DHCPv6 Local Server Overview on page 204
- DHCPv6 Relay Agent Overview on page 279
- Group-Specific DHCP Local Server Options on page 219
- Group-Specific DHCP Relay Options on page 292
- Guidelines for Configuring Interface Ranges on page 218

Guidelines for Configuring Interface Ranges

This topic describes guidelines to consider when configuring interface ranges for named interface groups for DHCP local server and DHCP relay. The guidelines refer to the following configuration statement:

user@host# set interface *interface-name* upto *upto-interface-name*

- The start subunit, **interface** ***interface-name*** , serves as the key for the stanza. The remaining configuration settings are considered attributes.
- If the subunit is not included, an implicit **.0** subunit is enforced. The implicit subunit is applied to all interfaces when autoconfiguration is enabled. For example, **interface ge-2/2/2** is treated as **interface ge-2/2/2.0**.
- Ranged entries contain the **upto** option, and the configuration applies to all interfaces within the specified range. The start of a ranged entry must be less than the end of the range. Discrete entries apply to a single interface, except in the case of autoconfiguration, in which a **0** (zero) subunit acts as a wildcard.
- Interface stanzas defined within the same router or switch context are dependent and can constrain each other—both DHCP local server and DHCP relay are considered. Interface stanzas defined across different router (switch) contexts are independent and do not constrain one another.
- Each interface stanza, whether discrete or ranged, has a unique start subunit across a given router context. For example, the following configuration is not allowed within the same group because **ge-1/0/0.10** is the start subunit for both.

  interface ge-1/0/0.10 upto ge-1/0/0.30
  interface ge-1/0/0.10

- Two groups cannot share interface space. For example, the following configuration is not allowed because the three stanzas share the same space and interfere with one another—interface **ge-1/0/0.26** is common to all three.

  dhcp-relay group diamond interface ge-1/0/0.10 upto ge-1/0/0.30
  dhcp-local-server group ruby interface ge-1/0/0.26
  dhcp-relay group sapphire interface ge-1/0/0.25 upto ge-1/0/0.35

- Two ranges cannot overlap, either within a group or across groups. Overlapping occurs when two interface ranges share common subunit space but neither range is a proper subset of the other. The following ranges overlap:

  interface ge-1/0/0.10 upto ge-1/0/0.30
  interface ge-1/0/0.20 upto ge-1/0/0.40

 Copyright © 2013, Juniper Networks, Inc.

- 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 217

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, and at the **[edit system services dhcp-local-server]** hierarchy level to set global DHCP local server configuration options. Statements configured at the **[edit system services dhcp-local-server group *group-name*]** hierarchy level apply only to the named group of interfaces, and override any global DHCP local server settings configured with the same statements at the **[edit system services dhcp-local-server]** hierarchy level.

DHCPv6 local server supports the same set of statements with the exception of the dynamic-profile statement.

- authentication—Configure the parameters the router sends to the external AAA server.
- dynamic-profile—Specify the dynamic profile that is attached to a group of interfaces.
- interface—Specify one or more interfaces, or a range of interfaces, that are within the specified group.
- overrides—Override the default configuration settings for the extended DHCP local server. For information, see "Overriding Default DHCP Local Server Configuration Settings" on page 220.

Related Documentation

- Grouping Interfaces with Common DHCP Configurations on page 217

Overriding Default DHCP Local Server Configuration Settings

Subscriber management enables you to override certain default DHCP and DHCPv6 local server configuration settings. You can override 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]** or **[edit system services dhcp-local-server dhcpv6]** 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*]** or **[edit system services dhcp-local-server dhcpv6 group]** 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]** or **[edit system services dhcp-local-server dhcpv6 group *group-name* interface]** hierarchy level.

To override default DHCP local server configuration settings:

1. Specify that you want to configure override options.

   Global override:

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

   Group level override:

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

   Per-interface override:

   ```
   [edit system services dhcp-local-server]
   user@host# edit group boston overrides interface fe-1/0/1.1
   ```

2. (Optional) Override the maximum number of DHCP clients allowed per interface.

   See "Specifying the Maximum Number of DHCP Clients Per Interface" on page 221.

3. (Optional) Configure DHCP client auto logout.

   See "Automatically Logging Out DHCP Clients" on page 226.

4. (Optional) Enable processing of information requests from clients.

   See "Enabling Processing of Client Information Requests" on page 227.

5. (Optional, DHCPv6 only) Specify a delegated pool name to use for DHCPv6 multiple address assignment.

   See "Specifying the Delegated Address Pool for IPv6 Prefix Assignment" on page 231.

6. (Optional, DHCPv6 only) Enable DHCPv6 rapid commit support.

Copyright © 2013, Juniper Networks, Inc.

See "Enabling DHCPv6 Rapid Commit Support" on page 228.

7. (Optional) Delete DHCP override settings.

   See "Deleting DHCP Local Server and DHCP Relay Override Settings" on page 229.

Related Documentation

- Group-Specific DHCP Local Server Options on page 219
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229

Specifying the Maximum Number of DHCP Clients Per Interface

By default, there is no limit to the number of DHCP local server or DHCP relay clients allowed on an interface. However, you can override the default setting and specify the maximum number of clients allowed per interface, in the range 1 through 500,000. When the number of clients on the interface reaches the specified limit, no additional DHCP Discover PDUs or DHCPv6 Solicit PDUs are accepted. When the number of clients subsequently drops below the limit, new clients are again accepted.

NOTE: The maximum number of DHCP (and DHCPv6) local server clients or DHCP (and DHCPv6) relay clients can also be specified by Juniper Networks VSA 26-143 during client login. The VSA-specified value always takes precedence if the **interface-client-limit** statement specifies a different number.

If the VSA-specified value differs with each client login, DHCP uses the largest limit set by the VSA until there are no clients on the interface.

To configure the maximum number of DHCP clients allowed per interface:

1. Specify that you want to configure override options.
   - For DHCP local server:

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

   - For DHCPv6 local server:

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

   - For DHCP relay agent:

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

   - For DHCPv6 relay agent:

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

2. Configure the maximum number of clients allowed per interface. (DHCP local server, DHCPv6 local server, DHCP relay agent and DHCPv6 relay agent all support the **interface-client-limit** statement.)

```
[edit system services dhcp-local-server overrides]
user@host# set interface-client-limit number
```

Related Documentation

- Overriding Default DHCP Local Server Configuration Settings on page 220
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229
- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276

DHCP Access, Access-Internal, and Destination Route Suppression

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

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

NOTE: In Junos OS releases prior to Release 13.2, the no-arp statement was used to override the default installation of host routes. In Junos OS Release 13.2 and later, the no-arp statement is deprecated and the function is replaced by the route-suppression statement.

Related Documentation

- Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default on page 223
- Extended DHCP Local Server Overview on page 200
- DHCPv6 Local Server Overview on page 204
- Extended DHCP Relay Agent Overview on page 276
- DHCPv6 Relay Agent Overview on page 279

Copyright © 2013, Juniper Networks, Inc.

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

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

NOTE: In Junos OS releases prior to Release 13.2, the **no-arp** statement was used to override the default installation of host routes. In Junos OS Release 13.2 and later, the **no-arp** statement is deprecated and the function is replaced by the **route-suppression** statement.

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

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

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

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

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

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

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

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

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

Related Documentation

- DHCP Access, Access-Internal, and Destination Route Suppression on page 222
- Extended DHCP Local Server Overview on page 200
- DHCPv6 Local Server Overview on page 204
- Extended DHCP Relay Agent Overview on page 276
- DHCPv6 Relay Agent Overview on page 279

Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server

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

Table 34 on page 224 indicates the action the router takes for DHCP local server snooped packets.

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

Table 34: Actions for DHCP Local Server Snooped Packets

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

To configure DHCP snooped packet forwarding for DHCP local server:

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

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

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

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

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

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

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

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

Related Documentation

- DHCP Snooping Support on page 299

Copyright © 2013, Juniper Networks, Inc.

DHCP Auto Logout Overview

This topic provides an introduction to the optional DHCP auto logout feature and includes the following sections:

- Auto Logout Overview on page 225
- How DHCP Identifies and Releases Clients on page 225
- Option 60 and Option 82 Requirements on page 226

Auto Logout Overview

Auto logout is an optional configuration for DHCP local server and DHCP relay agent that improves the efficiency of DHCP IP address assignment. Auto logout enables IP addresses to be immediately released and returned to the address pool when the addresses are no longer used by DHCP clients. 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 (on the routers), you might have an environment that includes set-top boxes (STB) that are often upgraded or replaced. Each time a STB is changed, the new STB repeats the DHCP discover process to obtain client configuration information and an IP address. DHCP views the new STB as a completely new client and assigns a new IP address—the previous IP address assigned to the client (the old STB) remains blocked and unavailable until the lease expires. If auto logout is configured in this situation, DHCP recognizes that the new STB is actually the same client and then immediately releases the original IP address. DHCP relay agent acts as a proxy client for auto logout and sends a DHCP release message to the DHCP server.

How DHCP Identifies and Releases Clients

The auto logout feature requires that DHCP explicitly identify clients. By default, DHCP local server and DHCP relay agent identify clients based on MAC address or Client Identifier. 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. The secondary identification method is based on the DHCP option 60 and option 82 information in DHCP discover messages.

Both the primary and secondary identification methods use subnet information to differentiate between clients. The primary identification method differentiates between two clients with the same MAC address (or same Client Identifier) if the clients are on different subnets. Similarly, the secondary identification method considers two clients

Copyright © 2013, Juniper Networks, Inc.

as different if they have the same option 60 and option 82 information, but different subnets.

DHCP local server and DHCP relay agent perform the following operations when auto logout is enabled and the secondary identification method identifies a duplicate client (that is, the discover packet is from an existing client).

- DHCP local server immediately releases the existing address.
- DHCP relay agent immediately releases the existing client and then sends a DHCP release packet to the DHCP server. Sending the release packet ensures that DHCP relay and the DHCP server are synchronized.

  If the DHCP relay receives a DISCOVER message from an existing client, the DHCP relay forwards the DISCOVER message to the DHCP server. The DHCP relay preserves the binding if the client's existing IP address is returned by the DHCP server. This behavior is not applicable if the proxy-mode override or client-discover-match functionality are enabled.

NOTE: **If the DHCP relay agent is in snoop mode, DHCP relay releases the client but does not send a release packet to the DHCP server if the discover packet is for a passive client (a client added as a result of snooped packets) or if the discover packet is a snooped packet.**

Option 60 and Option 82 Requirements

DHCP local server requires that the received discover packet include both DHCP option 60 and option 82. If either option is missing, DHCP local server cannot perform the secondary identification method and auto logout is not used.

DHCP relay agent requires that the received discover packet contain DHCP option 60. DHCP relay determines the option 82 value based on the guidelines provided in "DHCP Relay Agent Option 82 Value for Auto Logout" on page 310.

Related Documentation

- Automatically Logging Out DHCP Clients on page 226
- DHCP Relay Agent Option 82 Value for Auto Logout on page 310
- Clearing DHCP Bindings for Subscriber Access on page 234

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 that has the same DHCP option 60 and DHCP option 82 information as the existing client. DHCP then releases the existing client IP address without waiting for the normal lease expiration.

 Copyright © 2013, Juniper Networks, Inc.

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. (DHCP local server and DHCP relay agent both support the **client-discover-match** statement.)
   - For DHCP local server:

     [edit system services dhcp-local-server overrides]
     user@host# **set** client-discover-match

   - For DHCP relay:

     [edit forwarding-options dhcp-relay overrides]
     user@host# **set** client-discover-match

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 225
- DHCP Relay Agent Option 82 Value for Auto Logout on page 310
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229
- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276

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 169. 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" on page 172 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]
     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 220
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229
- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276

Enabling DHCPv6 Rapid Commit Support

You can configure the extended DHCPv6 local server to support the DHCPv6 Rapid Commit option (DHCPv6 option 14). When rapid commit is enabled on the extended DHCPv6 local server, 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

 Copyright © 2013, Juniper Networks, Inc.

a two-message exchange (Solicit and Reply) to configure clients, rather than the default four-method 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 override 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 220
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229
- Extended DHCP Local Server Overview on page 200

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 220
- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276

Use of DHCP Option 50 and DHCPv6 IA_NA Option to Request a Specific IP Address

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 200
- DHCPv6 Local Server Overview on page 204

Multiple Address Assignment for DHCPv6 Clients

Subscriber management (on the routers) or DHCP management (on the switches) enables you to assign multiple addresses to a single DHCPv6 client. Multiple address support is enabled by default, and is activated when the DHCPv6 local server receives a DHCPv6 Solicit message from a subscriber (or DHCP client) that contains multiple addresses.

For example, if you are implementing this feature on the routers, you might use the multiple address assignment feature in a networking environment in which a customer premises equipment (CPE) device requires a host address and a delegated prefix. In such an environment, you can configure subscriber management to assign both a DHCPv6 IA_NA (Identity Association for Non-Temporary Addresses) and an IA_PD (Identity Association for Prefix Delegation) address to the client (the CPE device).

- Multiple Address Assignment Using Local Address Pools or RADIUS on page 230
- Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment on page 231

Multiple Address Assignment Using Local Address Pools or RADIUS

You can use either local address pools or RADIUS when assigning multiple addresses to a DHCP client. When at least one address is successfully allocated, the router or switch creates a subscriber (or DHCP client) entry and binds the entry to the assigned address.

 Copyright © 2013, Juniper Networks, Inc.

If both addresses are successfully allocated, the router (or switch) creates a single subscriber (or DHCP client) entry and binds both addresses to that entry.

You can also configure a delegated address pool, which explicitly specifies the address pool that subscriber management (or DHCP management) uses to assign IPv6 prefixes for subscribers (or DHCP clients).

Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment

NOTE: EX Series switches do not support demux.

(On the routers only) Subscriber management provides a predefined variable that you can use to dynamically configure DHCPv6 multiple address assignment. You apply the Junos OS predefined variable, **$junos-subscriber-ipv6-multi-address**, as a demux source address in a dynamic profile. When the dynamic profile is attached to a subscriber, the variable is expanded to include both the host and prefix addresses. You use this variable instead of the **$junos-subscriber-ipv6-address** variable, which supports a single IPv6 address.

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.

Related Documentation

- Specifying the Delegated Address Pool for IPv6 Prefix Assignment on page 231
- Junos OS Predefined Variables on page 642

Specifying the Delegated Address Pool for IPv6 Prefix Assignment

You can explicitly specify a delegated address pool:

- On routers—Subscriber management uses the pool to assign IPv6 prefixes for subscribers. You can specify the delegated address pool globally, for a specific group of interfaces, or for a particular interface.
- On switches—DHCP management uses the pool to assign IPv6 prefixes for DHCP clients. You can specify the delegated address pool globally, for a specific group of interfaces, or for a particular interface.

NOTE: You can also use by Juniper Networks VSA 26-161 to specify the delegated address pool. The VSA-specified value always takes precedence over the **delegated-address** statement.

To configure the delegated address pool for DHCPv6 local server:

1. Specify that you want to configure override options.

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

2. Configure the delegated address pool.

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set delegated-pool paris-cable-12
```

Related Documentation

- Overriding Default DHCP Local Server Configuration Settings on page 220
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229
- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276

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 DPC 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.

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**

Related Documentation

- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 232
- Verifying and Managing the DHCP Maintain Subscribers Feature on page 233

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.

Copyright © 2013, Juniper Networks, Inc.

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 232
- Verifying and Managing the DHCP Maintain Subscribers Feature on page 233

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:

  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 232
- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 232

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     90:00:00:01:00:01   active   2011-10-17 11:38:47 PST
   192.168.32.3     90:00:00:02:00:01   active   2011-00-17 11:38:41 PST
   ```

2. Clear the binding you want to remove.

   ```
   user@host> clear dhcp server binding 192.168.32.1
   ```

3. Verify that the binding has been cleared.

   ```
   user@host> show dhcp server binding
   1 clients, (1 bound, 0 selecting, 0 renewing, 0 rebinding)

   IP address      Hardware address    Type     Lease expires at
   192.168.32.3    90:00:00:02:00:01   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
```


Copyright © 2013, Juniper Networks, Inc.

To clear bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

To clear all bindings over an interface. This example uses the wildcard option.

```
user@host> clear dhcp server binding ge-1/0/0. *
```

To clear bindings on top of a specific VLAN. This example clears all bindings on top of VLAN 100.

```
user@host> clear dhcp server binding ge-1/0/0:100
```

To clear bindings for a specific S-VLAN. This example clears bindings on S-VLAN 100-200.

```
user@host> clear dhcp server binding ge-1/0/0:100-200
```

To clear all bindings on top of all demux VLANs:

```
user@host> clear dhcp server binding demux0
```

To clear all bindings on top of an underlying interface. This example clears the bindings on all demux VLANS on top of interface **ae0**:

```
user@host> clear dhcp server binding ae0
```

To clear PPP bindings. This example uses the wildcard feature and clears the PPP bindings over interface **pp0.100** and **pp0.200**.

```
user@host> clear dhcp server binding pp0.*
```

To clear all bindings on an FPC. This example uses the wildcard feature and clears all DHCP bindings on FPC 1.

```
user@host> clear dhcp server binding ge-1/*
```

To clear all bindings on a PIC. This example uses the wildcard feature and clears all DHCP bindings on FPC 1, PIC 0.

```
user@host> clear dhcp server binding ge-1/0/*
```

To clear all bindings on a port. This example uses the wildcard feature and clears all DHCP bindings on FPC 1, PIC 0, port 0.

```
user@host> clear dhcp server binding ge-1/0/0.*
```

Related Documentation

- DHCP Auto Logout Overview on page 225
- Automatically Logging Out DHCP Clients on page 226

Configuring the Router to Strictly Enforce the Subscriber Scaling License

You can configure the router to strictly enforce the subscriber scaling feature, which is part of the Junos Subscriber Access Feature Pack license. The subscriber scaling feature specifies the maximum number of subscribers that can be logged in at any one time.

When you configure strict scaling license support, the router performs the following actions:

- Strictly enforces the subscriber scaling license and does not allow any grace period. When the number of logged-in subscriber reaches the number allowed by the scaling license, no additional subscribers are allowed to log in.
- Creates the informational log message, "90 percent of installed subscriber scale licenses in use" in **/var/log/messages**, to inform you when you have 10 percent of the total allowed licenses remaining. The router clears this condition when license usage falls below 90 percent. The log message is created again if the 90 percent usage is later reached.

To configure the router to strictly enforce the subscriber scaling license:

1. Specify that you want to configure subscriber management.

   [edit system services]
   user@host# **edit** subscriber-management

2. Configure the router to enforce the scaling license.

   [edit system services subscriber-management]
   user@host# **set** enforce-strict-scale-limit-license

Related Documentation

- Subscriber Access Licensing Overview on page 6
- For information about installing and managing Junos OS licenses, see the *Installation and Upgrade Guide*

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

This topic describes how to attach a dynamic profile to a DHCP subscriber interface or a DHCP client interface. When a DHCP subscriber or DHCP client logs in, the specified dynamic profile is instantiated and the services defined in the profile are applied to the interface.

This topic contains the following sections:

- Attaching a Dynamic Profile to All DHCP Subscriber or All DHCP Client Interfaces on page 236
- Attaching a Dynamic Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces on page 237

Attaching a Dynamic Profile to All DHCP Subscriber or All DHCP Client Interfaces

To attach a dynamic profile to all DHCP subscriber or all DHCP client interfaces:

1. At the DHCP configuration hierarchy, use the **dynamic-profile** statement to specify the name of the dynamic profile to attach to all interfaces.
   - For DHCP local server:

     [edit system services dhcp-local-server]

Copyright © 2013, Juniper Networks, Inc.

```
user@host# set dynamic-profile vod-profile-22
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# set dynamic-profile vod-profile-west
```

2. (Routers only) Optionally, you can configure the attribute to use when attaching the specified profile.

   You can include either the **aggregate-clients** option to enable multiple DHCP subscribers to share the same VLAN logical interface, or the **use-primary** option to specify that the primary dynamic profile is used. The **aggregate-clients** option does not apply to demux subscriber interfaces. The two options are mutually exclusive.

   - To enable multiple subscribers to share the same VLAN logical interface:

```
[edit system services dhcp-local-server dynamic-profile]
user@host# set aggregate-clients merge
```

   - To use the primary dynamic profile:

```
[edit forwarding-options dhcp-relay dynamic-profile]
user@host# set use-primary subscriber_profile
```

Attaching a Dynamic 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 217.

To attach a dynamic profile to a group of interfaces:

1. At the DHCP configuration hierarchy, specify the name of the interface group and the dynamic profile to attach to the group.

   - For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# set group boston dynamic-profile vod-profile-42
```

   - For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# set group quebec dynamic-profile vod-profile-east
```

2. (Routers only) Optionally, you can configure the attribute to use when attaching the specified profile.

   You can include either the **aggregate-clients** option to enable multiple DHCP subscribers to share the same VLAN logical interface, or the **use-primary** option to specify that the primary dynamic profile is used. The **aggregate-clients** option does not apply to demux subscriber interfaces. The two options are mutually exclusive.

   - To enable multiple subscribers to share the same VLAN logical interface:

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

```
user@host# set aggregate-clients merge
```

- To use the primary dynamic profile:

```
[edit forwarding-options dhcp-relay dynamic-profile]
user@host# set use-primary subscriber_profile
```

Related Documentation

- Dynamic Profiles Overview on page 638
- Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview on page 207
- Example: Configuring Dynamic Subscriber Interfaces on IP Demux Interfaces on page 781

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:

     ```
     [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 myPassworD1234
   ```

Related Documentation

- Extended DHCP Local Server Overview on page 200
- DHCPv6 Local Server Overview on page 204
- Extended DHCP Relay Agent Overview on page 276
- Using External AAA Authentication Services with DHCP on page 214
- For information about supported characters in passwords, see “Configuring Special Requirements for Plain-Text Passwords” in the *Junos OS Administration Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

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.

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-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***. (Not supported for DHCPv6 local server)
- **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 © 2013, 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 only)
- **relay-agent-remote-id**—The DHCPv6 Relay Agent Remote-ID option (option 37). (DHCPv6 local server only)
- **relay-agent-subscriber-id**—(On routers only) The DHCPv6 Relay Agent Subscriber-ID option (option 38). (DHCPv6 local server 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]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
     ```

2. Specify that you want to include optional information in the username. (DHCP local server, DHCPv6 local server, and DHCP relay agent all support the **username-include** statement.)

   ```
   [edit system services dhcp-local-server authentication]
   user@host# set username-include
   ```


Copyright © 2013, Juniper Networks, Inc.

3. (Optional) Specify the optional information you want to include in the username.

```
[edit system services dhcp-local-server authentication username-include]
user@host# set username-include circuit-type
user@host# set username-include domain-name isp55.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@isp55.com

Related Documentation

- Extended DHCP Local Server Overview on page 200
- DHCPv6 Local Server Overview on page 204
- Extended DHCP Relay Agent Overview on page 276
- Using External AAA Authentication Services with DHCP on page 214

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, the server sends a NAK message to the client to force it to begin the DHCP connection process again. Alternatively, the provider can wait for customers to make a service call about the network failures and then instruct them to power cycle their customer premises equipment to reinitiate the connection. Neither of these actions is timely or convenient for customers.

- On switches—Suppose you restructure the addressing scheme or change the server IP addresses that the DHCP server provides to clients. Without dynamic reconfiguration, the network typically clears the DHCP server binding table, but cannot inform the DHCP clients that their bindings have been cleared. Consequently, the DHCP client operates as though its IP address is still valid, but it is now unable to communicate over the access network, resulting in an outage. The DHCP local server needs to wait for the client to send a message to renew its lease or rebind to the server. In response, the server sends a NAK message to the client to force it to begin the DHCP connection process again. Alternatively, you can wait for users to notify you of the network failures and then instruct them to power cycle their equipment to reinitiate the connection. Neither of these actions is timely or convenient for users.

Dynamic Client/Server Interaction for DHCPv4

Dynamic reconfiguration for DHCPv4 is available through a partial implementation of RFC 3203, *DHCP Reconfigure Extension* for DHCPv4. It enables the DHCPv4 local server to send a message to the client to force reconfiguration.

The server sends a forcerenew message to a DHCPv4 client, initiating a message exchange. In response, DHCPv4 clients that support the forcerenew message then send a lease renewal message to the server. The server rejects the lease renewal request and sends a NAK to the client, causing the client to reinitiate the DHCP connection. A successful reconnection results in the reconfiguration of the DHCP client. Only the exchange of forcerenew, renew, and NAK messages is supported from RFC 3202. DHCP relay and DHCP relay proxy do not participate in the client reconfiguration or react to forcerenew messages other than to forward them to the client.

When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfiguring state and the local server sends a forcerenew message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber services or DHCP-managed services, such as forwarding and statistics, continue to work. Client statistics are not maintained in the interval between a successful reconfiguration and the subsequent client binding. When the server responds to the client renewal request with a NAK, the client entry is removed from the binding table and final statistics are reported. New statistics are collected when the client sends a discover message to establish a new session.

Dynamic Client/Server Interaction for DHCPv6

Dynamic reconfiguration for DHCPv6 is available through a partial implementation of RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*. It enables the DHCPv6 local server to send a message to the client to force reconfiguration.

DHCPv6 servers send reconfigure messages to DHCPv6 clients, initiating a message exchange. In response, DHCPv6 clients that support the reconfigure message transition to the renewing state and send a renew message to the server. The server returns a reply message with a lifetime of zero (0). The client transitions to the init state and sends a solicit message. The server sends an advertise message to indicate that it is available for service. The client sends a request for configuration parameters, which the server

 Copyright © 2013, Juniper Networks, Inc.

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.

Dynamic Configuration Options

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.

- To enable dynamic reconfiguration with default reconfiguration values for all DHCP clients, include the **reconfigure** statement at the **[edit system services dhcp-local-server]** hierarchy level for DHCPv4 clients, and at the **[edit system services dhcp-local-server dhcpv6]** hierarchy level for DHCPv6 clients.
- Alternatively, to enable dynamic reconfiguration for only the DHCP clients serviced by a specified group of interfaces, include the **reconfigure** statement at the **[edit system services dhcp-local-server group *group-name*]** hierarchy level for DHCPv4 clients, and at the **[edit system services dhcp-local-server dhcpv6 group *group-name*]** hierarchy level for DHCPv6 clients.

You can optionally modify the behavior of the reconfiguration process by including the appropriate 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.

Include the **attempts** statement to specify how many times the local server sends the forcerenew or reconfigure message to initiate client reconfiguration. Include the **timeout** statement to set the interval between the first and second attempts. The interval between each subsequent attempt doubles the previous value. 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.

By default, the DHCP client's original configuration is restored if all of the reconfiguration attempts fail. Include the **clear-on-abort** statement to delete the client instead.

You can configure an authentication token by including the **token** statement. 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.

In the event of a RADIUS-initiated disconnect (RID), the client is deleted by default. You can configure the client to be reconfigured instead of deleted by including the **radius-disconnect** statement. The client is deleted if all attempts to reconfigure the client fail.

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.

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.

Events that take place while a reconfiguration is in process take precedence over the reconfiguration. Table 35 on page 244 lists the actions taken in response to several different events.

Table 35: Action Taken for Events That Occur During a Reconfiguration

Event	Action
Server receives a discover (DHCPv4) or solicit (DHCPv6) message from the client.	Server drops packet and deletes client.
Server receives a request, renew, rebind, or init-reboot message from the client.	DHCPv4—Server sends NAK message and deletes client. DHCPv6—Server drops packet and deletes client. Server replies to renew message with lease time of zero (0).
Server receives a release or decline message from the client.	Server deletes client.
The client lease times out.	Server deletes client.
The **clear dhcp server binding** command is issued.	Server deletes client.
The **request dhcp server reconfigure** (DHCPv4) or **request dhcpv6 server reconfigure** (DHCPv6) command is issued.	Command is ignored.
GRES or DHCP restart occurs.	Reconfiguration process is halted.

Related Documentation

- Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 245

Copyright © 2013, Juniper Networks, Inc.

Configuring Extended DHCP Local Server Dynamic Client Reconfiguration

The DHCP local server can initiate reconfiguration of its clients to avoid extended outages because of server configuration changes. In addition to requesting that the DHCP local server initiate reconfiguration, you can specify the reconfiguration behavior.

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) Override the global configuration for a particular group of clients.

   For DHCPv4:

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

   For DHCPv6:

   ```
   [edit system services dhcp-local-server dhcpv6 group group-name]
   user@host# set reconfigure
   ```

3. (Optional) Configure how the server attempts reconfiguration.

   See "Configuring Dynamic Reconfiguration Attempts for DHCP Clients" on page 246.

4. (Optional) Configure the response to a failed reconfiguration.

   See "Configuring Deletion of the Client When Dynamic Reconfiguration Fails" on page 247.

5. (Optional) Configure the behavior in response to a RADIUS-initiated disconnect.

   See "Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect" on page 247.

6. (Optional) Configure a token for rudimentary server authentication.

   See "Configuring a Token for DHCP Local Server Authentication" on page 248.

7. (Optional) Initiate reconfiguration of some or all client bindings.

   See "Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings" on page 248.

8. (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 249.

Copyright © 2013, Juniper Networks, Inc.

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 messages. You can also specify how long the server waits between attempts. By default, eight attempts are made and the initial interval is two seconds.

Each successive attempt doubles the interval between attempts. For example, if the first value is 2, the first retry is attempted 2 seconds after the first attempt fails. The second retry is attempted 4 seconds after the first retry fails. The third retry is attempted 8 seconds after the second retry fails, and so on. A group configuration takes precedence over a DHCP local server configuration.

(Optional) To configure DHCP local server reconfiguration behavior for all DHCP clients:

1. Specify the number of reconfiguration attempts.

   For DHCPv4:

   ```
   [edit system services dhcp-local-server reconfigure]
   user@host# set attempts 5
   ```

   For DHCPv6:

   ```
   [edit system services dhcp-local-server dhcpv6 reconfigure]
   user@host# set attempts 5
   ```

2. Specify the interval between reconfiguration attempts.

   For DHCPv4:

   ```
   [edit system services dhcp-local-server reconfigure]
   user@host# set timeout 8
   ```

   For DHCPv6:

   ```
   [edit system services dhcp-local-server dhcpv6 reconfigure]
   user@host# set timeout 8
   ```

To override the global configuration for a particular group of clients, include the statements at the **[edit system services dhcp-local-server group *group-name* reconfigure]** hierarchy level or the **[edit system services dhcpv6 dhcp-local-server group *group-name* reconfigure]** hierarchy level.

Related Documentation

- Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- attempts on page 1462
- timeout on page 2024

 Copyright © 2013, Juniper Networks, Inc.

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 Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- clear-on-abort on page 1504

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 Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- radius-disconnect on page 1919
- trigger on page 2059

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 Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- token on page 2027

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
  ```


Copyright © 2013, Juniper Networks, Inc.

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:bd8:1111:2222::
  ```

- Specify the client ID of a DHCPv6 client.

  ```
  user@host> request dhcpv6 server reconfigure
    LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
  ```

- Specify the session ID of a DHCPv6 client.

  ```
  user@host> request dhcpv6 server reconfigure 5
  ```

- Specify the MAC address of a DHCPv4 client.

  ```
  user@host> request dhcp server reconfigure 12:23:34:45:56:67
  ```

- 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 Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- *request dhcp server reconfigure*

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.

(Optional) 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 particular 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 Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- **strict on page 2006**

DHCP Liveness Detection Overview

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

DHCP clients often do not send DHCP release messages prior to exiting the network. The discovery of their absence is dependent on existing DHCP lease time and release request mechanisms. These mechanisms are often considered insufficient when serving as session health checks for clients in a DHCP subscriber access or a DHCP-managed network. Because DHCP lease times are typically too long to provide an adequate response time for a session health failure, and configuring short DHCP lease times can pose an undue burden on control plane processing, implementing a DHCP liveness detection mechanism enables better monitoring of bound DHCP clients. When configured with a liveness detection protocol, if a given subscriber (or client) fails to respond to a configured number of consecutive liveness detection requests, the subscriber (or client) binding is deleted and its resources released.

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

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

- Provide more accurate time-based accounting of subscriber (or DHCP client) sessions.
- Better preserve router (switch) resources.

 Copyright © 2013, Juniper Networks, Inc.

- Help to reduce the window of vulnerability to some security attacks.

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

NOTE: **This release supports only BFD for DHCPv4 and DHCPv6 liveness detection.**

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

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

Related Documentation

- Configuring Detection of DHCP Local Server Client Connectivity on page 251
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 318

Configuring Detection of DHCP Local Server Client Connectivity

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

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

To configure liveness detection for DHCP local server:

1. Specify that you want to configure liveness detection.

   - For DHCP global configuration:

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

   - For DHCP group configuration:

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

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

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

   - For DHCP global configuration:

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

   - For DHCP group configuration:

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

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

   NOTE: In this release, the only method supported for liveness detection is Bidirectional Forwarding Detection (BFD).

   - For DHCP global configuration:

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

   - For DHCP group configuration:

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

4. Configure the liveness detection method as desired.

   See "Example: Configuring Group Liveness Detection for DHCP Local Server Clients" on page 269 for an example of how to configure DHCPv4 groups for DHCP local server liveness detection.

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

   - For DHCP global configuration:

Copyright © 2013, Juniper Networks, Inc.

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

- For DHCP group configuration:

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

Related Documentation

- DHCP Liveness Detection Overview on page 250
- Extended DHCP Local Server Overview on page 200
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 318
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

Verifying and Managing DHCP Local Server Configuration

Purpose View or clear information about client address bindings and statistics for the extended DHCP local server.

Action

- To display the address bindings in the client table on the extended DHCP local server:

```
user@host> show dhcp server binding
```

- To display extended DHCP local server statistics:

```
user@host> show dhcp server statistics
```

- To clear the binding state of a DHCP client from the client table on the extended DHCP local server:

```
user@host> clear dhcp server binding
```

- To clear all extended DHCP local server statistics:

```
user@host> clear dhcp server statistics
```

Related Documentation

- *Junos OS Operational Mode Commands*

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

- *Junos OS Operational Mode Commands*

Graceful Routing Engine Switchover

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

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

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

For more information about using graceful Routing Engine switchover, see the *Junos OS High Availability Library for Routing Devices*.

Related Documentation

- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276
- Subscriber Management Unified ISSU Support on page 7

 Copyright © 2013, Juniper Networks, Inc.

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 *Junos OS System Log Messages Reference*.)
- By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure global DHCP tracing operations.

- Specify tracing operations for DHCP local server and DHCP relay:

  ```
  [edit system processes dhcp-service]
  user@host# edit traceoptions
  ```

The tracing configuration is applied globally to all DHCP applications in every LS:RI. Configuration of event tracing on a per-LS:RI basis is not supported. DHCP tracing is configurable only in the default LS:RI. However, DHCP applications (local server or relay) do not have be configured in the default LS:RI. This behavior was different in software releases before Junos OS Release 11.4, where you had to configure a DHCP application in the default LS:RI in order to configure DHCP tracing, even when you wanted to run DHCP and trace its operations only in a nondefault LS:RI.

Copyright © 2013, Juniper Networks, Inc.

In the earlier software releases, you configured tracing statements at the **[edit system services dhcp-local-server]** and **[edit forwarding-options dhcp-relay]** hierarchy levels. These statements have been deprecated and hidden in favor of the statements at the **[edit system processes dhcp-service]** hierarchy level.

NOTE: **The deprecated statements may be removed from a future release; we recommend that you transition to the new statements.**

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 260.

The extended DHCP traceoptions operations are described in the following sections:

- Configuring the Extended DHCP Log Filename on page 256
- Configuring the Number and Size of Extended DHCP Log Files on page 257
- Configuring Access to the Extended DHCP Log File on page 257
- Configuring a Regular Expression for Extended DHCP Messages to Be Logged on page 258
- Configuring the Extended DHCP Tracing Flags on page 258
- Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged on page 259
- Tracing Extended DHCP Operations for Specific Interfaces on page 260

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.

 Copyright © 2013, Juniper Networks, Inc.

To change the filename:

- Specify a filename for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set file** ***filename***

- Specify a filename for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set file** ***filename***

Configuring the Number and Size of Extended DHCP Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format ***.number*.gz**. The newest archived file is **.0.gz** and the oldest archived file is ***.(maximum number)-1*.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

DHCP local server and DHCP relay agent both support the **files** and **size** options for the **traceoptions** statement and the **interface-traceoptions** statement. To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set file** ***filename*** **files** ***number*** **size** ***maximum-file-size***

- Specify the name, number, and size of the file used for the trace output for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set file** ***filename*** **files** ***number*** **size** ***maximum-file-size***

Configuring Access to the Extended DHCP Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

Copyright © 2013, Juniper Networks, Inc.

DHCP local server and DHCP relay agent both support the **world-readable** option and the **no-world-readable** option for the **traceoptions** statement and the **interface-traceoptions** statement. To specify that all users can read the log file:

- Configure the log file to be world-readable for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set file *filename* world-readable**

- Configure the log file to be world-readable for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set file *filename* world-readable**

To explicitly set the default behavior, in which the log file can only be read by the user who configured tracing:

- Configure the log file to be no-world-readable for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set file *filename* no-world-readable**

- Configure the log file to be no-world-readable for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set file *filename* no-world-readable**

Configuring a Regular Expression for Extended DHCP Messages to Be Logged

By default, the trace operation output includes all messages relevant to the logged events. You can refine the output by including regular expressions to be matched.

DHCP local server and DHCP relay agent both support the **match** option for the **traceoptions** statement and the **interface-traceoptions** statement. To configure regular expressions to be matched:

- Specify the regular expression for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set file *filename* match *regular-expression***

- Specify the regular expression for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set file *filename* match *regular-expression***

Configuring the Extended DHCP Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

DHCP local server and DHCP relay agent both support the **flag** option for the **traceoptions** statement and the **interface-traceoptions** statement. A smaller set of flags is supported for interface-level tracing than for global tracing. To configure the flags for the events to be logged:

- Specify the flags for global tracing operations.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit system processes dhcp-service traceoptions]
user@host# set flag flag
```

- Specify the flags for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]
user@host# set flag flag
```

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

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. A low severity level is less restrictive—filters out fewer messages—than a higher level. When you configure a severity level, all messages at that level and all higher (more restrictive) levels are logged.

The following list presents severity levels in order from lowest (least restrictive) to highest (most restrictive). This order also represents the significance of the messages; for example, **error** messages are of greater concern than **info** messages.

- **verbose**
- **info**
- **notice**
- **warning**
- **error**

The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all**. You can also specify **verbose** with the same result, because **verbose** is the lowest (least restrictive) severity level; it has nothing to do with the terseness or verbosity of the messages. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages. By default, the trace operation output includes only messages with a severity level of **error**.

DHCP local server and DHCP relay agent both support the **level** option for the **traceoptions** statement and the **interface-traceoptions** statement. To configure the flags for the events to be logged:

- Specify the severity level for global tracing operations.

```
[edit system processes dhcp-service traceoptions]
user@host# set level severity
```

- Specify the severity level for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]
user@host# set level severity
```

Tracing Extended DHCP Operations for Specific Interfaces

In addition to the global DHCP tracing operations, subscriber management enables you to trace extended DHCP operations for a specific interface or for a range of interfaces.

Configuring per-interface tracing is a two-step procedure. In the first step, you specify the tracing options that you want to use, such as file information and flags. In the second step, you enable the tracing operation on the specific interfaces.

To configure per-interface tracing operations:

1. Specify the tracing options you want to use.

   NOTE: **Per-interface tracing uses the same default tracing behavior as the global extended DHCP tracing operation. The default behavior is described in** "Tracing Extended DHCP Operations" on page 255.

   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 256.

      - Configure the number and size of the log files.

        See "Configuring the Number and Size of Extended DHCP Log Files" on page 257.

      - Configure access to the log file.

        See "Configuring Access to the Extended DHCP Log File" on page 257.

      - Configure a regular expression to filter logging events.

        See "Configuring a Regular Expression for Extended DHCP Messages to Be Logged" on page 258.

   c. (Optional) Specify tracing flag options.

      See "Configuring the Extended DHCP Tracing Flags" on page 258.

   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 259.

2. Enable tracing on an interface or interface range.

Copyright © 2013, Juniper Networks, Inc.

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 © 2013, Juniper Networks, Inc.

CHAPTER 8

DHCP Local Server Examples

- Example: Minimum Extended DHCP Local Server Configuration on page 263
- Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 263
- Example: Extended DHCPv6 Local Server Configuration on page 264
- Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine on page 265
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269

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.**

Related Documentation

- Extended DHCP Local Server Overview on page 200

Example: Extended DHCP Local Server Configuration with Optional Pool Matching

This example shows an extended DHCP local server configuration that includes optional IPv4 address-assignment pool matching and interface groups. For pool matching, this configuration specifies that the DHCP local server first check the response from an external authentication authority (for example, RADIUS) and use the Framed-IPv6-Pool

attribute to determine the address-assignment pool to use for the client address. If no external authority match is found, the DHCP local server then uses ip-address-first matching together with the option 82 information to match the named address range for client IPv4 address assignment. The option 82 matching must also be included in the address-assignment pool configuration.

```
[edit system services]
dhcp-local-server {
  group group_one {
    interface fe-0/0/2.0;
    interface fe-0/0/2.1;
  }
  group group_two {
    interface fe-0/0/3.0;
    interface fe-0/0/3.1;
  }
  pool-match-order {
    external-authority
    ip-address-first;
    option-82;
  }
}
```

NOTE: The interface type in this topic is just an example. The fe- interface type is not supported by EX Series switches.

Related Documentation

- Extended DHCP Local Server Overview on page 200
- Address-Assignment Pools Overview on page 167

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 v679M8vt;
      username-include {
        user-prefix wallybrown;
        domain-name isp55.com;
      }
    }
    group group_two {
      authentication {
        password P$55qw4$$;
        username-include {
          user-prefix south5;
          domain-name isp55.com;
```

 Copyright © 2013, Juniper Networks, Inc.

```
            }
          }
          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 &tIUEI*7688+;
  }
}
profile isp-bos-metro-fiber-basic {
  accounting-order radius;
  authentication-order radius;
  radius {
    authentication-server 192.168.1.250;
    accounting-server 192.168.1.250;
  }
  accounting {
    order radius;
    accounting-stop-on-failure;
    accounting-stop-on-access-deny;
    update-interval 10;
    statistics time;
  }
}
```

Related Documentation

- DHCPv6 Local Server Overview on page 204

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 265
- Overview on page 266
- Configuration on page 266
- Verification on page 268

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 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
```


Copyright © 2013, Juniper Networks, Inc.

```
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
```

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.

```
[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 268

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
```

 Copyright © 2013, Juniper Networks, Inc.

```
Filter: RE-protect
Counters:
Name                                                Bytes              Packets
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 208
- Dynamic Firewall Filters Overview on page 1132
- *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*
- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276

Example: Configuring Group Liveness Detection for DHCP Local Server Clients

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

- Requirements on page 269
- Overview on page 270
- Configuration on page 270

Requirements

- Juniper Networks MX Series routers
- Juniper Networks EX Series switches
- Junos OS Release 12.1R1 or later
- Junos OS Release 12.3R2 or later for EX Series switches
- Configure DHCP local server. See "Extended DHCP Local Server Overview" on page 200.

Overview

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

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

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

Configuration

Step-by-Step Procedure

To configure group liveness detection for DHCP local server:

1. Specify that you want to configure liveness detection.

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

2. Specify that you want to configure liveness detection for a specific DHCP local server group.

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

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

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

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

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

5. Configure the detection time threshold (in milliseconds) at which a trap is produced.

   [edit system services dhcp-local-server group local_group_1 liveness-detection method bfd]
   user@host# set detection-time threshold 30000

6. Configure the time (in milliseconds) for which BFD holds a session up notification.

   [edit system services dhcp-local-server group local_group_1 liveness-detection method bfd]

Copyright © 2013, Juniper Networks, Inc.

```
user@host# set holddown-interval 50
```

7. Configure the BFD minimum transmit and receive interval (in milliseconds).

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

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

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

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

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

9. Configure a multiplier value for the detection time.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method bfd]
user@host# set multiplier 100
```

10. Disable the ability for BFD interval timers to change or adapt to network situations.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method bfd]
user@host# set no-adaptation
```

11. Configure the BFD session mode.

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

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

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

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

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

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

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

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

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

```
[edit]
regress@montag# show system
services {
  dhcp-local-server {
    group local_group_1 {
      liveness-detection {
        failure-action clear-binding-if-interface-up;
        method {
          bfd {
            version automatic;
            minimum-interval 45000;
            minimum-receive-interval 60000;
            multiplier 100;
            no-adaptation;
            transmit-interval {
              minimum-interval 45000;
              threshold 60000;
            }
            detection-time {
              threshold 30000;
            }
            session-mode automatic;
            holddown-interval 50;
          }
        }
      }
    }
  }
}
```

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

Related Documentation

- Extended DHCP Local Server Overview on page 200
- DHCP Liveness Detection Overview on page 250
- Configuring Detection of DHCP Local Server Client Connectivity on page 251

 Copyright © 2013, Juniper Networks, Inc.

PART 4

DHCP Relay Agent for Subscriber Access

- DHCP Relay Agent Overview on page 275
- Configuring DHCP Relay Agent on page 285
- DHCP Relay Agent Examples on page 341

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 9

DHCP Relay Agent Overview

- Extended DHCP Relay Agent Overview on page 276
- DHCPv6 Relay Agent Overview on page 279
- DHCP Relay Proxy Overview on page 280
- Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview on page 282

Extended DHCP Relay Agent Overview

Copyright © 2013, Juniper Networks, Inc.

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. For more information about how to use the DHCP relay agent in a video/IPTV application, see the *Junos OS Feature Guides*.

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 the *Routing Policy Feature Guide for Routing Devices*.

You can also configure the extended DHCP relay agent to support IPv6 clients. See "DHCPv6 Relay Agent Overview" on page 279 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. See the **"[edit forwarding-options dhcp-relay] Hierarchy Level" on page 1398** for the complete DHCP relay agent syntax.

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:

- Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers on page 278
- DHCP Liveness Detection on page 279

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 © 2013, 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 279
- Access and Access-Internal Routes for Subscriber Management on page 687
- Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview on page 207
- Using External AAA Authentication Services with DHCP on page 214
- DHCP Relay Proxy Overview on page 280
- Graceful Routing Engine Switchover on page 254
- Subscriber Management Unified ISSU Support on page 7
- Verifying and Managing DHCP Relay Configuration on page 332
- Tracing Extended DHCP Operations on page 255
- Example: Minimum DHCP Relay Agent Configuration on page 341
- Example: DHCP Relay Agent Configuration with Multiple Clients and Servers on page 342
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 349
- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 353
- Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine on page 265

DHCPv6 Relay Agent Overview

The DHCPv6 relay agent enhances the extended 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 extended DHCP local server and the extended DHCP relay agent, and can be enabled on the same interface as either the extended 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]**

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214
- Grouping Interfaces with Common DHCP Configurations on page 217
- Group-Specific DHCP Relay Options on page 292
- Overriding the Default DHCP Relay Configuration Settings on page 293
- Configuring Passwords for Usernames on page 238
- Creating Unique Usernames for DHCP Clients on page 239
- Verifying and Managing DHCPv6 Local Server Configuration on page 253
- Example: Extended DHCPv6 Local Server Configuration on page 264

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

 Copyright © 2013, Juniper Networks, Inc.

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: **Extended DHCP relay proxy is not supported for the J Series Services Routers DHCP server. Also, 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.

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.

Copyright © 2013, Juniper Networks, Inc.

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 276
- Enabling DHCP Relay Proxy Mode on page 329
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 318

Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview

The router's DHCP support enables you to attach a dynamic profile to a DHCP subscriber interface. When a DHCP subscriber logs in, the router instantiates the specified dynamic profile and then applies the services defined in the profile to the interface.

You can attach dynamic profiles to all interfaces or you can specify a particular group of interfaces to which the profile is attached. Both the DHCP local server and the DHCP relay agent support the attachment of dynamic profiles to interfaces.

You can enable the following optional features when the dynamic profile is attached. The two options cannot be used together.

- Enable multiple DHCP subscribers to share the same VLAN logical interface. The firewall filters, CoS schedulers, and IGMP configuration of the clients are merged.
- Specify the primary dynamic profile that is instantiated when the first subscriber logs in.

Multiple DHCP Subscribers Sharing the Same VLAN Logical Interface

The **aggregate-clients** statement specifies that the router merge the firewall filters, CoS schedulers, and IGMP configuration of multiple DHCP clients that are on the same VLAN logical interface (for example, multiple clients belonging to the same household). You

 Copyright © 2013, Juniper Networks, Inc.

can configure the aggregate-clients support for all interfaces or for a group of interfaces. The **aggregate-clients** statement provides the option of either merging (chaining) or replacing software components for each client.

By default, the feature is disabled and a single DHCP client is allowed per VLAN when a dynamic profile is associated with the VLAN logical interface.

When you specify the **merge** option, the router aggregates the software components for multiple subscribers as follows:

- Firewall filters—The filters are chained together using the precedence as the order of execution. If the same firewall filter is attached multiple times, the filter is executed only once.
- CoS schedulers—The different CoS schedulers are merged as if the scheduler map has multiple schedulers. The merge operation for the individual traffic-control-profiles parameters (shaping-rate, delay-buffer-rate, guaranteed-rate) preserves the maximum value for each parameter.
- IGMP configuration—The current IGMP configuration is replaced with the configuration of the newest DHCP client.

When you specify the **replace** option, the entire logical interface is replaced whenever a new client logs in to the network using the same VLAN logical interface. For example, if a customer subscribes to voice, video, and data services on the network, when a voice client logs in, instead of applying a specific voice filter for only that service, the entire voice, video, and data filter chain is applied.

NOTE: **You cannot use a dynamic demux interface to represent multiple subscribers in a dynamic profile attached to an interface. One dynamic demux interface represents one subscriber. Do not configure the aggregate-clients option when attaching a dynamic profile to a demux interface for DHCP.**

Primary Dynamic Profile

The **use-primary** option enables you to specify the primary dynamic profile that 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.

This feature can conserve logical interfaces in a network where dynamic IP demux interfaces are used to represent subscribers. To conserve interfaces, make sure the primary profile that you specify does not create a demux interface, but provides the initial policies for the primary interface subscriber.

Related Documentation

- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 10

Configuring DHCP Relay Agent

- DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 286
- Guidelines for Configuring Support for DHCP Duplicate Clients on page 287
- Configuring DHCP Duplicate Client Support on page 288
- Using External AAA Authentication Services with DHCP on page 288
- Grouping Interfaces with Common DHCP Configurations on page 290
- Guidelines for Configuring Interface Ranges on page 291
- Group-Specific DHCP Relay Options on page 292
- Overriding the Default DHCP Relay Configuration Settings on page 293
- Overwriting giaddr Information on page 294
- Replacing the DHCP Relay Request and Release Packet Source Address on page 295
- Overriding Option 82 Information on page 295
- Using Layer 2 Unicast Transmission for DHCP Packets on page 296
- Trusting Option 82 Information on page 296
- DHCP Access, Access-Internal, and Destination Route Suppression on page 297
- Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default on page 298
- Specifying the Maximum Number of DHCP Clients Per Interface on page 298
- DHCP Snooping Support on page 299
- Configuring DHCP Snooping for DHCP Relay Agent on page 301
- Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 301
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 306
- DHCP Auto Logout Overview on page 308
- DHCP Relay Agent Option 82 Value for Auto Logout on page 310
- Automatically Logging Out DHCP Clients on page 311
- Sending Release Messages When Clients Are Deleted on page 312
- Subscriber Binding Retention During Interface Delete Events on page 313

Copyright © 2013, Juniper Networks, Inc.

- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 313
- Verifying and Managing the DHCP Maintain Subscribers Feature on page 314
- Clearing DHCP Bindings for Subscriber Access on page 315
- DHCP Liveness Detection Overview on page 316
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 318
- Disabling DHCP Relay on page 319
- Disabling Automatic Binding of Stray DHCP Requests on page 320
- DHCP Options and Selective Traffic Processing Overview on page 321
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 323
- Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings on page 325
- Enabling and Disabling Insertion of Option 82 Information on page 325
- Configuring Server Groups on page 328
- Configuring Active Server Groups on page 328
- Enabling DHCP Relay Proxy Mode on page 329
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 329
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 330
- Verifying and Managing DHCP Relay Configuration on page 332
- Verifying and Managing DHCPv6 Relay Configuration on page 333
- Tracing Extended DHCP Operations on page 333
- Disabling DHCP Relay on page 339

DHCP Duplicate Client Differentiation Using Client Subinterface Overview

In some network environments, client IDs and MAC addresses might not be unique, resulting in duplicate clients. For example, two network adapters might be manufactured with the same hardware address, resulting in a duplicate MAC address among the DHCP clients attached to the router (or switch). A duplicate DHCP client occurs when a client attempts to get a lease, and that client has the same client ID or the same MAC address as an existing DHCP client.

When DHCP server receives a request from a new client that has a duplicate ID or MAC address, DHCP server terminates the address lease for the existing client and returns the address to its original address pool. DHCP server then assigns a new address and lease to the new client.

By default, both DHCP local server and DHCP relay use the subnet information to differentiate between duplicate clients. However, in some cases, this level of differentiation is not adequate. For example, when multiple subinterfaces share the same underlying

 Copyright © 2013, Juniper Networks, Inc.

loopback interface with the same preferred source address, the interfaces appear to be on the same subnet. In this situation, the default configuration prevents duplicate clients.

You can provide greater differentiation between duplicate clients by configuring DHCP to consider the client subinterface when duplicate clients occur. In this optional configuration, DHCP uniquely identifies:

- The subnet on which the client resides
- The subinterface on which the client resides
- The client within the subnet

Related Documentation

- Configuring DHCP Duplicate Client Support on page 214
- Guidelines for Configuring Support for DHCP Duplicate Clients on page 213

Guidelines for Configuring Support for DHCP Duplicate Clients

This topic describes the guidelines for configuring DHCP to include the client subinterface in order to distinguish between duplicate clients (clients with the same MAC address or client ID) in a subscriber access or in a DHCP-managed environment.

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

- The optional DHCP duplicate client support feature is used for DHCPv4 clients. For DHCPv6, client identification is independent of MAC address.
- For DHCP relay agent configuration:
  - DHCP relay must be configured to insert option 82, regardless of whether or not the incoming packet has option 82.
  - Option 82 must include the Agent Circuit ID suboption (suboption 1).
  - Option 82 must be the interface name, not the interface description.
  - DHCP server must echo option 82 in the server's reply. This is required because of the following:
    - The giaddr inserted by DHCP relay is the same for duplicate clients on different subinterfaces. The DHCP local server uses option 82 when allocating the IP address.
    - DHCP relay uses the echoed option 82 to learn the client subinterface and to construct the client key.
- For the Layer 3 wholesale model:
  - The wholesaler and retailer logical system/routing instances must have the same **duplicate-clients-on-interface** statement configuration.
  - For DHCP relay, the wholesaler and the retailer routing contexts must both be configured with the Agent Circuit ID suboption (suboption 1) in option 82.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 212
- Configuring DHCP Duplicate Client Support on page 214

Configuring DHCP Duplicate Client Support

You can optionally configure DHCP local server and DHCP relay to include a client subinterface when distinguishing between two clients that have the same MAC address or client ID. The configuration is a global setting for each logical system/routing instance.

To configure DHCP local server to include the client subinterface:

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

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

2. Configure the optional duplicate client support.

   [edit system services dhcp-local-server]
   user@host# **set** duplicate-clients-on-interface

To configure DHCP relay agent to include the client subinterface:

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

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

2. Configure the optional duplicate client support.

   [edit system services dhcp-relay]
   user@host# **set** duplicate-clients-on-interface

Related Documentation

- DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 212
- Guidelines for Configuring Support for DHCP Duplicate Clients on page 213

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.

Copyright © 2013, Juniper Networks, Inc.

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.
   - 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 238.

3. (Optional) Configure optional features to create a unique username.

   See "Creating Unique Usernames for DHCP Clients" on page 239.

Related Documentation

- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276
- DHCPv6 Local Server Overview on page 204
- DHCPv6 Relay Agent Overview on page 279

Copyright © 2013, Juniper Networks, Inc.

Grouping Interfaces with Common DHCP Configurations

You use the group feature to group together a set of interfaces and then apply a common DHCP configuration to the named interface group. The extended DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent all support interface groups.

The following steps create a DHCP local server group; the steps are similar for the DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent.

To configure a DHCP local server interface group:

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

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

2. Create the group and assign a name.

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

3. Specify the names of one or more interfaces on which the extended DHCP application is enabled. You can repeat the **interface** ***interface-name*** statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

   ```
   [edit system services dhcp-local-server group boston]
   user@host# set interface fe-1/0/1.1
   user@host# set interface fe-1/0/1.2
   ```

4. (Optional) You can use the **upto** option to specify a range of interfaces for a group.

   ```
   [edit system services dhcp-local-server group boston]
   user@host# set interface fe-1/0/1.3 upto fe-1/0/1.9
   ```

5. (Optional) You can use the **exclude** option to exclude a specific interface or a specified range of interfaces from the group. For example:

   ```
   [edit system services dhcp-local-server group boston]
   user@host# set interface fe-1/0/1.1 upto fe-1/0/1.102
   user@host# set interface fe-1/0/1.6 exclude
   user@host# set interface fe-1/0/1.70 upto fe-1/0/1.80 exclude
   ```

Related Documentation

- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276
- DHCPv6 Local Server Overview on page 204
- DHCPv6 Relay Agent Overview on page 279
- Group-Specific DHCP Local Server Options on page 219
- Group-Specific DHCP Relay Options on page 292
- Guidelines for Configuring Interface Ranges on page 218

Copyright © 2013, Juniper Networks, Inc.

Guidelines for Configuring Interface Ranges

This topic describes guidelines to consider when configuring interface ranges for named interface groups for DHCP local server and DHCP relay. The guidelines refer to the following configuration statement:

user@host# set interface *interface-name* upto *upto-interface-name*

- The start subunit, **interface** ***interface-name*** , serves as the key for the stanza. The remaining configuration settings are considered attributes.
- If the subunit is not included, an implicit **.0** subunit is enforced. The implicit subunit is applied to all interfaces when autoconfiguration is enabled. For example, **interface ge-2/2/2** is treated as **interface ge-2/2/2.0**.
- Ranged entries contain the **upto** option, and the configuration applies to all interfaces within the specified range. The start of a ranged entry must be less than the end of the range. Discrete entries apply to a single interface, except in the case of autoconfiguration, in which a **0** (zero) subunit acts as a wildcard.
- Interface stanzas defined within the same router or switch context are dependent and can constrain each other—both DHCP local server and DHCP relay are considered. Interface stanzas defined across different router (switch) contexts are independent and do not constrain one another.
- Each interface stanza, whether discrete or ranged, has a unique start subunit across a given router context. For example, the following configuration is not allowed within the same group because **ge-1/0/0.10** is the start subunit for both.

  ```
  interface ge-1/0/0.10 upto ge-1/0/0.30
  interface ge-1/0/0.10
  ```

- Two groups cannot share interface space. For example, the following configuration is not allowed because the three stanzas share the same space and interfere with one another—interface **ge-1/0/0.26** is common to all three.

  ```
  dhcp-relay group diamond interface ge-1/0/0.10 upto ge-1/0/0.30
  dhcp-local-server group ruby interface ge-1/0/0.26
  dhcp-relay group sapphire interface ge-1/0/0.25 upto ge-1/0/0.35
  ```

- Two ranges cannot overlap, either within a group or across groups. Overlapping occurs when two interface ranges share common subunit space but neither range is a proper subset of the other. The following ranges overlap:

  ```
  interface ge-1/0/0.10 upto ge-1/0/0.30
  interface ge-1/0/0.20 upto ge-1/0/0.40
  ```

- A range can contain multiple nested ranges. A nested range is a proper subset of another range. When ranges are nested, the smallest matching range applies.

  In the following example, the three ranges nest properly:

  ```
  interface ge-1/0/0.10 upto ge-1/0/0.30
  interface ge-1/0/0.12 upto ge-1/0/0.15 exclude
  interface ge-1/0/0.25 upto ge-1/0/0.29 exclude
  ```

Copyright © 2013, Juniper Networks, Inc.

- 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 217

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" on page 328.
- **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" on page 250.
- **overrides**—Override the default configuration settings for the extended DHCP relay agent. For information, see "Overriding the Default DHCP Relay Configuration Settings" on page 293.
- **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-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 323.
- **relay-option-82**—(DHCPv4 only) Enable or disable the insertion of option 82 information in packets destined for a DHCP server. For information, see "Enabling and Disabling Insertion of Option 82 Information" on page 325.
- **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

 Copyright © 2013, Juniper Networks, Inc.

by a RADIUS server or a provisioning server. For more information, see "Default Subscriber Service Overview" on page 1127.

Related Documentation

- Grouping Interfaces with Common DHCP Configurations on page 217

Overriding the Default DHCP Relay Configuration Settings

You can override the default DHCP and DHCPv6 relay agent 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]** hierarchy level.
- To configure overrides for DHCPv6 relay, use the supported statements at the **[edit forwarding-options dhcp-relay dhcpv6]** hierarchy level.

To override default DHCP relay agent configuration settings:

1. Specify that you want to configure override options.

   Global override:

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

   Group-level override:

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

   Per-interface override:

   ```
   [edit forwarding-options dhcp-relay]
   user@host# edit group boston interface fe-1/0/1.2 overrides
   ```

2. (DHCPv4 only) Enable DHCP relay proxy mode.

   See "Enabling DHCP Relay Proxy Mode" on page 329.

3. (DHCPv4 only) Overwrite the giaddr in DHCP packets that the DHCP relay agent forwards.

   See "Overwriting giaddr Information" on page 294.

4. (DHCPv4 only) Replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).

Copyright © 2013, Juniper Networks, Inc.

See "Replacing the DHCP Relay Request and Release Packet Source Address" on page 295.

5. (DHCPv4 only) Override the DHCP relay agent information option (option 82) in DHCP packets.

   See "Overriding Option 82 Information" on page 295.

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 296.

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 296.

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 221.

9. (DHCPv4 only) Configure client auto logout.

   See "DHCP Auto Logout Overview" on page 225.

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" on page 301.

11. (DHCPv4 and DHCPv6) Send release messages to the DHCP server when clients are deleted.

    See "Sending Release Messages When Clients Are Deleted" on page 312.

12. (DHCPv4 only) Disable the DHCP relay agent on specific interfaces.

    See "Disabling DHCP Relay" on page 319.

13. (DHCPv4 and DHCPv6) Disable automatic binding of stray DHCP requests.

    See "Disabling Automatic Binding of Stray DHCP Requests" on page 320.

Related Documentation

- Group-Specific DHCP Relay Options on page 292
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229

Overwriting giaddr Information

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.

Copyright © 2013, Juniper Networks, Inc.

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 276
- Overriding the Default DHCP Relay Configuration Settings on page 293

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.

   ```
   [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 276
- Overriding the Default DHCP Relay Configuration Settings on page 293

Overriding Option 82 Information

You can configure the DHCP relay agent to add or remove the DHCP relay agent information option (option 82) in DHCP packets.

This feature causes the DHCP relay agent to perform one of the following actions, depending on the configuration:

- If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.
- If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.

Copyright © 2013, Juniper Networks, Inc.

To override the default option 82 information in DHCP packets destined for a DHCP server:

1. Specify that you want to configure override options.

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

2. Specify that the option 82 information in DHCP packets is overwritten.

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

Related Documentation

- Extended DHCP Relay Agent Overview on page 276
- Overriding the Default DHCP Relay Configuration Settings on page 293

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 276
- Overriding the Default DHCP Relay Configuration Settings on page 293

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 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

Copyright © 2013, Juniper Networks, Inc.

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 276
- Overriding the Default DHCP Relay Configuration Settings on page 293

DHCP Access, Access-Internal, and Destination Route Suppression

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

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

NOTE: In Junos OS releases prior to Release 13.2, the no-arp statement was used to override the default installation of host routes. In Junos OS Release 13.2 and later, the no-arp statement is deprecated and the function is replaced by the route-suppression statement.

Related Documentation

- Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default on page 223
- Extended DHCP Local Server Overview on page 200
- DHCPv6 Local Server Overview on page 204
- Extended DHCP Relay Agent Overview on page 276
- DHCPv6 Relay Agent Overview on page 279

Copyright © 2013, Juniper Networks, Inc.

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

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

NOTE: In Junos OS releases prior to Release 13.2, the **no-arp** statement was used to override the default installation of host routes. In Junos OS Release 13.2 and later, the **no-arp** statement is deprecated and the function is replaced by the **route-suppression** statement.

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

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

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

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

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

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

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

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

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

Related Documentation

- DHCP Access, Access-Internal, and Destination Route Suppression on page 222
- Extended DHCP Local Server Overview on page 200
- DHCPv6 Local Server Overview on page 204
- Extended DHCP Relay Agent Overview on page 276
- DHCPv6 Relay Agent Overview on page 279

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

Copyright © 2013, Juniper Networks, Inc.

Discover PDUs or DHCPv6 Solicit PDUs are accepted. When the number of clients subsequently drops below the limit, new clients are again accepted.

NOTE: The maximum number of DHCP (and DHCPv6) local server clients or DHCP (and DHCPv6) relay clients can also be specified by Juniper Networks VSA 26-143 during client login. The VSA-specified value always takes precedence if the **interface-client-limit** statement specifies a different number.

If the VSA-specified value differs with each client login, DHCP uses the largest limit set by the VSA until there are no clients on the interface.

To configure the maximum number of DHCP clients allowed per interface:

1. Specify that you want to configure override options.
   - For DHCP local server:

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

   - For DHCPv6 local server:

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

   - For DHCP relay agent:

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

   - For DHCPv6 relay agent:

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

2. Configure the maximum number of clients allowed per interface. (DHCP local server, DHCPv6 local server, DHCP relay agent and DHCPv6 relay agent all support the **interface-client-limit** statement.)

   [edit system services dhcp-local-server overrides]
   user@host# **set** interface-client-limit *number*

Related Documentation

- Overriding Default DHCP Local Server Configuration Settings on page 220
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229
- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276

DHCP Snooping Support

DHCP snooping provides DHCP security on the router by filtering incoming messages. When DHCP snooping is enabled, the router differentiates between trusted and untrusted

interfaces, and forwards messages from trusted sources while rejecting the untrusted messages.

In Junos OS, DHCP snooping is enabled in a routing instance when you configure either the **dhcp-relay** statement at the **[edit forwarding-options]** hierarchy level, or the **dhcp-local-server** statement at the **[edit system services]** hierarchy level in that routing instance. However, depending on the Junos OS release, the router processes the snooped packets differently, as described in the following list:

- In Junos OS Release 10.0 and earlier, the router processes snooped packets normally.
- In Junos OS Release 10.1 and later, he router discards snooped packets by default. To enable normal processing of snooped packets in Junos OS Release 10.1 and later, you must explicitly configure the **allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level.

You can configure DHCP snooping support for the following:

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

  In a separate procedure, you can set a global configuration to specify whether the DHCPv4 relay agent forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces. The router also uses the global DHCP relay agent snooping configuration to determine whether to forward or drop snooped BOOTREPLY packets.
- DHCPv6 relay agent—As you can with snooping support for the DHCPv4 relay agent, you can override the default DHCPv6 relay agent snooping configuration on the router to explicitly enable or disable snooping support globally, for a named group of interfaces, or for a specific interface with a named group of interfaces.

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

  Unlike the DHCPv4 relay agent, the DHCPv6 relay agent does not support global configuration of forwarding support for DHCPv6 snooped packets.
- DHCP local server—Configure whether DHCP local server forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces.

Related Documentation

- Configuring DHCP Snooping for DHCP Relay Agent on page 301
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 223
- Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 343

Copyright © 2013, Juniper Networks, Inc.

Configuring DHCP Snooping for DHCP Relay Agent

DHCP relay agent uses a two-part configuration to determine how to handle DHCP snooped packets. First, you enable or disable snooping support for DHCP relay agent and, optionally, override the default snooping configuration. Then you configure the forwarding action for snooped clients, which specifies whether DHCP relay agent forwards or drops snooped traffic.

To configure DHCP snooping for DHCP relay agent:

1. (DHCPv4 and DHCPv6) Enable or disable DHCP snooping. You can configure DHCP snooping globally, for a named group of interfaces, or for a specific interface.

   See "Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent" on page 301.

2. (DHCPv4 only) Configure snooped packets forwarding support.

   See "Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent" on page 306.

Related Documentation

- DHCP Snooping Support on page 299
- Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 301
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 306
- Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 343

Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent

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

The second procedure, which applies only to DHCPv4 relay agent, is described in "Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent" on page 306, and configures the **forward-snooped-clients** statement, which determines whether the snooped packets are forwarded or dropped, depending on the type of interface.

NOTE: The router has a default global setting that specifies whether DHCP snooping support is enabled or disabled for DHCP relay. In Junos OS Release 10.0 and earlier, DHCP snooping is enabled by default. In Junos OS Release 10.1 and later, DHCP snooping is disabled by default.

You can override the default global DHCP snooping configuration and explicitly enable or disable DHCP snooping support. You can configure the explicit snooping support globally, for a group of interfaces, or for a specific interface in a group.

- To enable DHCP relay or DHCPv6 relay snooping support, include the **allow-snooped-clients** option in the **overrides** statement.
- To disable DHCP relay or DHCPv6 relay snooping support, include the **no-allow-snooped-clients** option in the **overrides** statement.

To enable or disable DHCP snooping support globally:

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

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

   - For DHCPv6 relay agent:

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

2. Specify that you want to override the default configuration.
   - For DHCP relay agent:

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

   - For DHCPv6 relay agent:

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

3. Enable or disable DHCP snooping support.
   - To enable DHCP snooping:
     - For DHCP relay agent:

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

     - For DHCPv6 relay agent:

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

   - To disable DHCP snooping:
     - For DHCP relay agent:

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

     - For DHCPv6 relay agent:

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

Copyright © 2013, Juniper Networks, Inc.

For example, to enable global DHCP snooping support :

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

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

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

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

   - For DHCPv6 relay agent:

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

2. Specify the named group.
   - For DHCP relay agent:

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

   - For DHCPv6 relay agent:

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

3. Specify that you want to override the default configuration.
   - For DHCP relay agent:

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

   - For DHCPv6 relay agent:

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

4. Enable or disable DHCP snooping support.
   - To enable DHCP snooping:
     - For DHCP relay agent:

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

     - For DHCPv6 relay agent:

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

   - To disable DHCP snooping:

- For DHCP relay agent:

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

- For DHCPv6 relay agent:

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

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

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

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

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

   - For DHCP relay agent:

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

   - For DHCPv6 relay agent:

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

2. Specify the named group containing the interface.

   - For DHCP relay agent:

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

   - For DHCPv6 relay agent:

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

3. Specify the interface for which you want to configure DHCP snooping.

   - For DHCP relay agent:

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

   - For DHCPv6 relay agent:

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

4. Specify that you want to override the default configuration on the interface.

   - For DHCP relay agent:

Copyright © 2013, Juniper Networks, Inc.

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

- For DHCPv6 relay agent:

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

5. Enable or disable DHCP snooping support.

   - To enable DHCP snooping:

     - For DHCP relay agent:

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

     - For DHCPv6 relay agent:

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

   - To disable DHCP snooping:

     - For DHCP relay agent:

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

     - For DHCPv6 relay agent:

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

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

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

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

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

```
                allow-snooped-clients;
            }
        }
    }
  }
 }
}
```

Related Documentation

- Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 306
- DHCP Snooping Support on page 299
- Overriding the Default DHCP Relay Configuration Settings on page 293

Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

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

DHCP relay uses a two-part configuration to determine how to handle DHCP snooped packets. This topic describes how you use the **forward-snooped-clients** statement to manage whether DHCP relay agent forwards or drops snooped packets, depending on the type of interface on which the packets are snooped. In the other part of the DHCP relay agent snooping configuration, which is described in "Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent" on page 301, you enable or disable the DHCP relay snooping feature.

Table 36 on page 307 shows the action the router or switch takes on snooped packets when DHCP snooping is enabled by the **allow-snooped-clients** statement. Table 37 on page 307 shows the action the router (or switch) takes on snooped packets when DHCP snooping is disabled by the **no-allow-snooped-clients** statement.

The router or switch also uses the configuration of the DHCP relay agent forwarding support to determine how to handle snooped BOOTREPLY packets. Table 38 on page 307 shows the action the router (or switch) takes for the snooped BOOTREPLY packets.

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

Copyright © 2013, Juniper Networks, Inc.

Table 36: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled

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

Table 37: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Disabled

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

Table 38: Actions for Snooped BOOTREPLY Packets

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

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

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

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

2. Enable DHCP snooped packet forwarding.

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

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

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

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

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

Related Documentation

- DHCP Snooping Support on page 299
- Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 301

DHCP Auto Logout Overview

This topic provides an introduction to the optional DHCP auto logout feature and includes the following sections:

- Auto Logout Overview on page 308
- How DHCP Identifies and Releases Clients on page 309
- Option 60 and Option 82 Requirements on page 309

Auto Logout Overview

Auto logout is an optional configuration for DHCP local server and DHCP relay agent that improves the efficiency of DHCP IP address assignment. Auto logout enables IP addresses to be immediately released and returned to the address pool when the addresses are no longer used by DHCP clients. 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 (on the routers), you might have an environment that includes set-top boxes (STB) that are often upgraded or replaced. Each time a STB is changed, the new STB repeats the DHCP discover process to obtain client configuration information and an IP address. DHCP views the new STB as a completely new client and assigns a new IP address—the previous IP address assigned to the client (the old STB) remains blocked and unavailable until the lease expires. If auto logout is configured in this situation, DHCP recognizes that the new STB is actually the same client and then immediately releases the original IP address. DHCP relay agent acts as a proxy client for auto logout and sends a DHCP release message to the DHCP server.

 Copyright © 2013, Juniper Networks, Inc.

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. 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. The secondary identification method is based on the DHCP option 60 and option 82 information in DHCP discover messages.

Both the primary and secondary identification methods use subnet information to differentiate between clients. The primary identification method differentiates between two clients with the same MAC address (or same Client Identifier) if the clients are on different subnets. Similarly, the secondary identification method considers two clients as different if they have the same option 60 and option 82 information, but different subnets.

DHCP local server and DHCP relay agent perform the following operations when auto logout is enabled and the secondary identification method identifies a duplicate client (that is, the discover packet is from an existing client).

- DHCP local server immediately releases the existing address.
- DHCP relay agent immediately releases the existing client and then sends a DHCP release packet to the DHCP server. Sending the release packet ensures that DHCP relay and the DHCP server are synchronized.

  If the DHCP relay receives a DISCOVER message from an existing client, the DHCP relay forwards the DISCOVER message to the DHCP server. The DHCP relay preserves the binding if the client's existing IP address is returned by the DHCP server. This behavior is not applicable if the proxy-mode override or client-discover-match functionality are enabled.

NOTE: **If the DHCP relay agent is in snoop mode, DHCP relay releases the client but does not send a release packet to the DHCP server if the discover packet is for a passive client (a client added as a result of snooped packets) or if the discover packet is a snooped packet.**

Option 60 and Option 82 Requirements

DHCP local server requires that the received discover packet include both DHCP option 60 and option 82. If either option is missing, DHCP local server cannot perform the secondary identification method and auto logout is not used.

Copyright © 2013, Juniper Networks, Inc.

DHCP relay agent requires that the received discover packet contain DHCP option 60. DHCP relay determines the option 82 value based on the guidelines provided in "DHCP Relay Agent Option 82 Value for Auto Logout" on page 310.

Related Documentation

- Automatically Logging Out DHCP Clients on page 226
- DHCP Relay Agent Option 82 Value for Auto Logout on page 310
- Clearing DHCP Bindings for Subscriber Access on page 234

DHCP Relay Agent Option 82 Value for Auto Logout

Table 39 on page 310 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 39: 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
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

Copyright © 2013, Juniper Networks, Inc.

Table 39: DHCP Relay Agent Option 82 Value for Auto Logout *(continued)*

DHCP Relay Agent Configuration Settings				giaddr in non-snooped packet	Action Taken
DHCP Relay Configured with Option 82	Discover Packet Contains Option 82	Override "trust-option-82"	Override "always-write-option-82"		
Yes	Yes	Yes	Yes	–	Overwrite the configured option 82

Related Documentation

- DHCP Auto Logout Overview on page 225
- Automatically Logging Out DHCP Clients on page 226

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 that has the same DHCP option 60 and DHCP option 82 information as the 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. (DHCP local server and DHCP relay agent both support the **client-discover-match** statement.)
   - For DHCP local server:

     [edit system services dhcp-local-server overrides]
     user@host# **set** client-discover-match

   - For DHCP relay:

     [edit forwarding-options dhcp-relay overrides]

```
user@host# set client-discover-match
```

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 225
- DHCP Relay Agent Option 82 Value for Auto Logout on page 310
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229
- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276

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: In Junos OS Release 10.1 and earlier, DHCP relay sends a release message to the DHCP server when the **client-discover-match** statement is included as a DHCP relay override. In Junos OS Release 10.2 and later, 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 © 2013, Juniper Networks, Inc.

```
user@host# set send-release-on-delete
```

Related Documentation

- Extended DHCP Relay Agent Overview on page 276
- Overriding the Default DHCP Relay Configuration Settings on page 293

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 DPC 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.

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**

Related Documentation

- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 232
- Verifying and Managing the DHCP Maintain Subscribers Feature on page 233

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.

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 232
- Verifying and Managing the DHCP Maintain Subscribers Feature on page 233

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:

  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 232
- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 232

 Copyright © 2013, Juniper Networks, Inc.

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     90:00:00:01:00:01   active   2011-10-17 11:38:47 PST
   192.168.32.3     90:00:00:02:00:01   active   2011-00-17 11:38:41 PST
   ```

2. Clear the binding you want to remove.

   ```
   user@host> clear dhcp server binding 192.168.32.1
   ```

3. Verify that the binding has been cleared.

   ```
   user@host> show dhcp server binding
   1 clients, (1 bound, 0 selecting, 0 renewing, 0 rebinding)

   IP address      Hardware address    Type     Lease expires at
   192.168.32.3    90:00:00:02:00:01   active   2011-00-17 11:38:41 PST
   ```

The following examples show variations of the clear DHCP binding feature. The examples use the DHCP local server version of the commands.

NOTE: **IP demux interfaces are not supported by the show and clear DHCP bindings commands for DHCP local server and DHCP relay agent.**

To clear all bindings:

user@host> **clear dhcp server binding all**

To clear bindings on a specific interface:

user@host> **clear dhcp server binding interface fe-0/0/2**

To clear all bindings over an interface. This example uses the wildcard option.

```
user@host> clear dhcp server binding ge-1/0/0. *
```

To clear bindings on top of a specific VLAN. This example clears all bindings on top of VLAN 100.

```
user@host> clear dhcp server binding ge-1/0/0:100
```

To clear bindings for a specific S-VLAN. This example clears bindings on S-VLAN 100-200.

```
user@host> clear dhcp server binding ge-1/0/0:100-200
```

To clear all bindings on top of all demux VLANs:

```
user@host> clear dhcp server binding demux0
```

To clear all bindings on top of an underlying interface. This example clears the bindings on all demux VLANS on top of interface **ae0**:

```
user@host> clear dhcp server binding ae0
```

To clear PPP bindings. This example uses the wildcard feature and clears the PPP bindings over interface **pp0.100** and **pp0.200**.

```
user@host> clear dhcp server binding pp0.*
```

To clear all bindings on an FPC. This example uses the wildcard feature and clears all DHCP bindings on FPC 1.

```
user@host> clear dhcp server binding ge-1/*
```

To clear all bindings on a PIC. This example uses the wildcard feature and clears all DHCP bindings on FPC 1, PIC 0.

```
user@host> clear dhcp server binding ge-1/0/*
```

To clear all bindings on a port. This example uses the wildcard feature and clears all DHCP bindings on FPC 1, PIC 0, port 0.

```
user@host> clear dhcp server binding ge-1/0/0.*
```

Related Documentation

- DHCP Auto Logout Overview on page 225
- Automatically Logging Out DHCP Clients on page 226

DHCP Liveness Detection Overview

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

DHCP clients often do not send DHCP release messages prior to exiting the network. The discovery of their absence is dependent on existing DHCP lease time and release request mechanisms. These mechanisms are often considered insufficient when serving as session health checks for clients in a DHCP subscriber access or a DHCP-managed network. Because DHCP lease times are typically too long to provide an adequate

 Copyright © 2013, Juniper Networks, Inc.

response time for a session health failure, and configuring short DHCP lease times can pose an undue burden on control plane processing, implementing a DHCP liveness detection mechanism enables better monitoring of bound DHCP clients. When configured with a liveness detection protocol, if a given subscriber (or client) fails to respond to a configured number of consecutive liveness detection requests, the subscriber (or client) binding is deleted and its resources released.

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

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

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

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

NOTE: **This release supports only BFD for DHCPv4 and DHCPv6 liveness detection.**

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

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

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring Detection of DHCP Local Server Client Connectivity on page 251
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 318

Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity

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

NOTE: **You can also configure DHCP liveness detection for DHCP local server.**

To configure liveness detection for DHCP relay:

1. Specify that you want to configure liveness detection.
   - For DHCP global configuration:

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

   - For DHCP group configuration:

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

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

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

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

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

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

   - For DHCP group configuration:

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

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

Copyright © 2013, Juniper Networks, Inc.

NOTE: In this release, the only method supported for liveness detection is Bidirectional Forwarding Detection (BFD).

- For DHCP global configuration:

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

- For DHCP group configuration:

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

5. Configure the liveness detection method as desired.

   See "Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients" on page 357 for an example of how to globally configure DHCP relay liveness detection.

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

   - For DHCP global configuration:

     [edit forwarding-options dhcp-relay liveness-detection]
     user@host# edit failure-action *action*

   - For DHCP group configuration:

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

Related Documentation

- Extended DHCP Relay Agent Overview on page 276
- DHCP Liveness Detection Overview on page 250
- Configuring Detection of DHCP Local Server Client Connectivity on page 251
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

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 276
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229

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: In Junos OS Release 10.4 and later, 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. In those releases, DHCP relay drops stray requests and forwards a NAK to the client when authentication is configured. Otherwise, DHCP relay attempts to bind the requesting client. In those releases, DHCP relay proxy always drops stray requests and forwards a NAK to the client, regardless of the authentication configuration.

- 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
   ```


Copyright © 2013, Juniper Networks, Inc.

2. Specify that you want to configure overrides.

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

3. Disable automatic binding for the group.

   [edit forwarding-options dhcp-relay group boston overrides]
   user@host# set no-bind-on-request

To disable automatic binding of stray requests on a specific interface:

1. Specify the named group of which the interface is a member.

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

2. Specify the interface on which you want to disable automatic binding.

   [edit forwarding-options dhcp-relay group boston]
   user@host# edit interface fe-1/0/1.2

3. Specify that you want to configure overrides.

   [edit forwarding-options dhcp-relay group boston interface fe-1/0/1.2]
   user@host# edit overrides

4. Disable automatic binding on the interface.

   [edit forwarding-options dhcp-relay group boston interface fe-1/0/1.2 overrides]
   user@host# set no-bind-on-request

Related Documentation

- Extended DHCP Relay Agent Overview on page 276
- Overriding the Default DHCP Relay Configuration Settings on page 293

DHCP Options and Selective Traffic Processing Overview

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

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

DHCP relay agent to differentiate between the two servers and to take the correct action for the subscriber.

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

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

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

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

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

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

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

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

Use the following match criteria to filter client traffic:

- **equals**—Your specified string is an exact match to the option string in client traffic.
- **starts-with**—Your specified string is a subset of the option string in client traffic, starting with the left-most character. For example, your configuration of the string "test" is a subset of "test123" in the client's option string, and matches the **starts-with** criteria.

Copyright © 2013, Juniper Networks, Inc.

- default-action—The option string in client traffic does not satisfy any match criteria, or no match criteria are configured.

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

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

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

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

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

Related Documentation

- Extended DHCP Relay Agent Overview on page 276
- Grouping Interfaces with Common DHCP Configurations on page 217
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 323
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 349
- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 353

Using DHCP Option Information to Selectively Process DHCP Client Traffic

The following procedure describes the steps you use to configure DHCP relay agent to selectively process client traffic—selective processing uses DHCP or DHCPv6 option information to identify, filter, and process client traffic. To configure DHCPv6 support you use the procedure at the **[edit forwarding-options dhcp-relay dhcpv6]** hierarchy level.

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

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

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

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

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

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

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

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

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

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

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

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

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

   For example, to select traffic that has an option 60 (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
   ```

Related Documentation

- DHCP Options and Selective Traffic Processing Overview on page 321
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 349
- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 353

Copyright © 2013, Juniper Networks, Inc.

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 on page 321
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 323
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 349

Enabling and Disabling Insertion of Option 82 Information

You can enable or disable support for the DHCP relay agent information option (option 82) in packets destined for a DHCP server. You can configure option 82 support globally or for a named group of interfaces.

To restore the default behavior (option 82 information is not inserted into DHCP packets), you use the **delete relay-option-82** statement.

To configure support for the DHCP relay agent information option 82, you use the **relay-option-82** statement.

The following sections describe the option 82 operations you can configure:

- Configuring Agent Circuit ID Information on page 325
- Configuring an Option 82 Prefix on page 326
- Using a Textual Description in Option 82 on page 327

Configuring Agent Circuit ID Information

You use the **relay-option-82** statement to enable insertion of option 82 information in DHCP packets. You must also specify at least the **circuit-id** statement to include the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option.

If you specify the **circuit-id** statement, the format of the Agent Circuit ID information for Fast Ethernet (**fe**) or Gigabit Ethernet (**ge**) interfaces is one of the following, depending on your network configuration:

- For Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual local area networks (VLANs) or stacked VLANs (S-VLANs):

  (fe | ge)-*fpc/pic/port*

- For Fast Ethernet or Gigabit Ethernet interfaces that use VLANs:

(fe | ge)-*fpc/pic/port*:*vlan-id*

- For Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs:

  (fe | ge)-*fpc/pic/port*:*svlan-id-vlan-id*

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. Specify insertion of the Agent Circuit ID suboption.

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

Configuring an Option 82 Prefix

You can include an optional prefix to the base option 82 information in DHCP packets destined for a DHCP server.

The prefix is separated from the option 82 Agent Circuit ID information by a colon (:), and 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 router configured with the **host-name** statement at the **[edit system]** hierarchy level.
- If you include the **logical-system-name** option, the DHCP relay agent uses the logical system name configured with the **logical-system** statement at the **[edit logical-system]** hierarchy level.
- If you include the **routing-instance-name** option, the DHCP relay agent uses the routing instance name configured with the **routing-instance** statement at the **[edit routing-instances]** hierarchy level or at the **[edit logical-system *logical-system-name* routing-instances]** hierarchy level.

If you include the hostname and either or both of the logical system name and the routing instance name in the prefix, the hostname is followed by a forward slash (/). If you include both the logical system name and the routing instance name in the prefix, these values are separated by a semicolon (;).

The following examples show several possible formats for the Agent Circuit ID 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:

Copyright © 2013, Juniper Networks, Inc.

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 Agent Circuit ID format. For Fast Ethernet or Gigabit Ethernet interfaces that do not use VLANs or S-VLANs, neither the ***vlan-id*** value nor the ***svlan-id*** value appears.

To configure an optional 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. Specify insertion of the Agent Circuit ID information.

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

3. Specify that the prefix is included in the option 82 information. In this example, the prefix includes the hostname and logical system name

   [edit forwarding-options dhcp-relay relay-option-82 circuit-id]
   user@host# **set** prefix **host-name logical-system-name**

Using a Textual Description in Option 82

By default, when DHCP option 82 is inserted into client packets, the Agent Circuit ID suboption includes the interface identifier. You can optionally configure that the Agent Circuit ID suboption 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 Agent Circuit ID suboption for static interfaces. The textual description is configured 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 interface identifier.

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. Specify insertion of the Agent Circuit ID information.

   [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

Configuring 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 276

Configuring Active 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.

Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

Copyright © 2013, Juniper Networks, Inc.

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 276
- Grouping Interfaces with Common DHCP Configurations on page 217

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 280
- Overriding the Default DHCP Relay Configuration Settings on page 293

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.

- 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 use the DHCPv4 Option 82 Agent Circuit ID suboption (suboption 1). This configuration is useful in a dual-stack environment, which has both DHCPv4 and DHCPv6 subscribers that reside over the same underlying logical interface. The router checks for the option 82 suboption 1 value and inserts it into the outgoing packets. If no DHCPv4 binding exists or if the binding does not have an option 82 suboption 1 value, the router sends the packets without adding an option 18.

NOTE: **If you specify one of the optional configurations, and the specified information does not exist (for example, there is no interface description), DHCPv6 relay ignores the optional configuration and inserts the default interface identifier in the packets.**

To insert the DHCPv6 Interface-ID option (option 18) in DHCPv6 packets:

1. Configure the DHCPv6 relay to include option 18.

   [edit forwarding-options dhcp-relay dhcpv6]
   user@host# **edit** relay-agent-interface-id

2. (Optional) Specify the prefix to include in option 18.

   [edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
   user@host# **set** prefix *prefix*

3. (Optional) Specify that option 18 include the textual description of the interface. You can specify either the **logical** interface description or the **device** interface description.

   [edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
   user@host# **set** use-interface-description **(logical | device)**

4. (Optional) Specify that option 18 use the DHCPv4 Option 82 Agent Circuit ID suboption (suboption 1) value.

   [edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
   user@host# **set** use-option-82

Related Documentation

- Extended DHCP Relay Agent Overview on page 276

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

This topic describes how to attach a dynamic profile to a DHCP subscriber interface or a DHCP client interface. When a DHCP subscriber or DHCP client logs in, the specified

Copyright © 2013, Juniper Networks, Inc.

dynamic profile is instantiated and the services defined in the profile are applied to the interface.

This topic contains the following sections:

- Attaching a Dynamic Profile to All DHCP Subscriber or All DHCP Client Interfaces on page 331
- Attaching a Dynamic Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces on page 331

Attaching a Dynamic Profile to All DHCP Subscriber or All DHCP Client Interfaces

To attach a dynamic profile to all DHCP subscriber or all DHCP client interfaces:

1. At the DHCP configuration hierarchy, use the **dynamic-profile** statement to specify the name of the dynamic profile to attach to all interfaces.
   - For DHCP local server:

     [edit system services dhcp-local-server]
     user@host# **set** dynamic-profile **vod-profile-22**

   - For DHCP relay agent:

     [edit forwarding-options dhcp-relay]
     user@host# **set** dynamic-profile **vod-profile-west**

2. (Routers only) Optionally, you can configure the attribute to use when attaching the specified profile.

   You can include either the **aggregate-clients** option to enable multiple DHCP subscribers to share the same VLAN logical interface, or the **use-primary** option to specify that the primary dynamic profile is used. The **aggregate-clients** option does not apply to demux subscriber interfaces. The two options are mutually exclusive.

   - To enable multiple subscribers to share the same VLAN logical interface:

     [edit system services dhcp-local-server dynamic-profile]
     user@host# **set aggregate-clients merge**

   - To use the primary dynamic profile:

     [edit forwarding-options dhcp-relay dynamic-profile]
     user@host# **set use-primary subscriber_profile**

Attaching a Dynamic 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 217.

To attach a dynamic profile to a group of interfaces:

1. At the DHCP configuration hierarchy, specify the name of the interface group and the dynamic profile to attach to the group.

- For DHCP local server:

  [edit system services dhcp-local-server]
  user@host# **set group boston** dynamic-profile **vod-profile-42**

- For DHCP relay agent:

  [edit forwarding-options dhcp-relay]
  user@host# **set group quebec** dynamic-profile **vod-profile-east**

2. (Routers only) Optionally, you can configure the attribute to use when attaching the specified profile.

   You can include either the **aggregate-clients** option to enable multiple DHCP subscribers to share the same VLAN logical interface, or the **use-primary** option to specify that the primary dynamic profile is used. The **aggregate-clients** option does not apply to demux subscriber interfaces. The two options are mutually exclusive.

   - To enable multiple subscribers to share the same VLAN logical interface:

     [edit system services dhcp-local-server dynamic-profile]
     user@host# **set aggregate-clients merge**

   - To use the primary dynamic profile:

     [edit forwarding-options dhcp-relay dynamic-profile]
     user@host# **set use-primary subscriber_profile**

Related Documentation

- Dynamic Profiles Overview on page 638
- Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview on page 207
- Example: Configuring Dynamic Subscriber Interfaces on IP Demux Interfaces on page 781

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**

- To display extended DHCP relay agent statistics:

  user@host> **show dhcp relay statistics**

- To clear the binding state of DHCP relay agent clients:

  user@host> **clear dhcp relay binding**

- To clear all extended DHCP relay agent statistics:

  user@host> **clear dhcp relay statistics**

Related Documentation

- *Junos OS Operational Mode Commands*

 Copyright © 2013, Juniper Networks, Inc.

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

- *Junos OS Operational Mode Commands*

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 *Junos OS System Log Messages Reference*.)

- By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure global DHCP tracing operations.

- Specify tracing operations for DHCP local server and DHCP relay:

  [edit system processes dhcp-service]
  user@host# **edit** traceoptions

The tracing configuration is applied globally to all DHCP applications in every LS:RI. Configuration of event tracing on a per-LS:RI basis is not supported. DHCP tracing is configurable only in the default LS:RI. However, DHCP applications (local server or relay) do not have be configured in the default LS:RI. This behavior was different in software releases before Junos OS Release 11.4, where you had to configure a DHCP application in the default LS:RI in order to configure DHCP tracing, even when you wanted to run DHCP and trace its operations only in a nondefault LS:RI.

In the earlier software releases, you configured tracing statements at the **[edit system services dhcp-local-server]** and **[edit forwarding-options dhcp-relay]** hierarchy levels. These statements have been deprecated and hidden in favor of the statements at the **[edit system processes dhcp-service]** hierarchy level.

NOTE: **The deprecated statements may be removed from a future release; we recommend that you transition to the new statements.**

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.

Copyright © 2013, Juniper Networks, Inc.

For information about configuring per-interface tracing options, see "Tracing Extended DHCP Operations for Specific Interfaces" on page 260.

The extended DHCP traceoptions operations are described in the following sections:

- Configuring the Extended DHCP Log Filename on page 335
- Configuring the Number and Size of Extended DHCP Log Files on page 335
- Configuring Access to the Extended DHCP Log File on page 336
- Configuring a Regular Expression for Extended DHCP Messages to Be Logged on page 336
- Configuring the Extended DHCP Tracing Flags on page 337
- Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged on page 337
- Tracing Extended DHCP Operations for Specific Interfaces on page 338

Configuring the Extended DHCP Log Filename

By default, the name of the file that records trace output is **jdhcpd**. You can specify a different name by including the **file** option. DHCP local server and DHCP relay agent both support the **file** option for the **traceoptions** statement and the **interface-traceoptions** statement.

To change the filename:

- Specify a filename for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set file** ***filename***

- Specify a filename for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set file** ***filename***

Configuring the Number and Size of Extended DHCP Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format ***.number*.gz**. The newest archived file is **.0.gz** and the oldest archived file is ***.(maximum number)-1*.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats

until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

DHCP local server and DHCP relay agent both support the **files** and **size** options for the **traceoptions** statement and the **interface-traceoptions** statement. To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set file *filename* files *number* size *maximum-file-size***

- Specify the name, number, and size of the file used for the trace output for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set file *filename* files *number* size *maximum-file-size***

Configuring Access to the Extended DHCP Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

DHCP local server and DHCP relay agent both support the **world-readable** option and the **no-world-readable** option for the **traceoptions** statement and the **interface-traceoptions** statement. To specify that all users can read the log file:

- Configure the log file to be world-readable for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set file *filename* world-readable**

- Configure the log file to be world-readable for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set file *filename* world-readable**

To explicitly set the default behavior, in which the log file can only be read by the user who configured tracing:

- Configure the log file to be no-world-readable for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set file *filename* no-world-readable**

- Configure the log file to be no-world-readable for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set file *filename* no-world-readable**

Configuring a Regular Expression for Extended DHCP Messages to Be Logged

By default, the trace operation output includes all messages relevant to the logged events. You can refine the output by including regular expressions to be matched.

 Copyright © 2013, Juniper Networks, Inc.

DHCP local server and DHCP relay agent both support the **match** option for the **traceoptions** statement and the **interface-traceoptions** statement. To configure regular expressions to be matched:

- Specify the regular expression for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set file** ***filename*** **match** ***regular-expression***

- Specify the regular expression for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set file** ***filename*** **match** ***regular-expression***

Configuring the Extended DHCP Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

DHCP local server and DHCP relay agent both support the **flag** option for the **traceoptions** statement and the **interface-traceoptions** statement. A smaller set of flags is supported for interface-level tracing than for global tracing. To configure the flags for the events to be logged:

- Specify the flags for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# **set flag** ***flag***

- Specify the flags for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# **set flag** ***flag***

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

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. A low severity level is less restrictive—filters out fewer messages—than a higher level. When you configure a severity level, all messages at that level and all higher (more restrictive) levels are logged.

The following list presents severity levels in order from lowest (least restrictive) to highest (most restrictive). This order also represents the significance of the messages; for example, **error** messages are of greater concern than **info** messages.

- **verbose**
- **info**
- **notice**
- **warning**
- **error**

The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all**. You can also specify **verbose** with the same result, because **verbose** is the lowest (least restrictive) severity level; it has nothing to do with the terseness or verbosity of the messages. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages. By default, the trace operation output includes only messages with a severity level of **error**.

DHCP local server and DHCP relay agent both support the **level** option for the **traceoptions** statement and the **interface-traceoptions** statement. To configure the flags for the events to be logged:

- Specify the severity level for global tracing operations.

  [edit system processes dhcp-service traceoptions]
  user@host# set level *severity*

- Specify the severity level for per-interface tracing operations.

  [edit system processes dhcp-service interface-traceoptions]
  user@host# set level *severity*

Tracing Extended DHCP Operations for Specific Interfaces

In addition to the global DHCP tracing operations, subscriber management enables you to trace extended DHCP operations for a specific interface or for a range of interfaces.

Configuring per-interface tracing is a two-step procedure. In the first step, you specify the tracing options that you want to use, such as file information and flags. In the second step, you enable the tracing operation on the specific interfaces.

To configure per-interface tracing operations:

1. Specify the tracing options you want to use.

   NOTE: Per-interface tracing uses the same default tracing behavior as the global extended DHCP tracing operation. The default behavior is described in "Tracing Extended DHCP Operations" on page 255.

   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.

Copyright © 2013, Juniper Networks, Inc.

- Configure the name for the file used for the trace output.

  See "Configuring the Extended DHCP Log Filename" on page 256.

- Configure the number and size of the log files.

  See "Configuring the Number and Size of Extended DHCP Log Files" on page 257.

- Configure access to the log file.

  See "Configuring Access to the Extended DHCP Log File" on page 257.

- Configure a regular expression to filter logging events.

  See "Configuring a Regular Expression for Extended DHCP Messages to Be Logged" on page 258.

c. (Optional) Specify tracing flag options.

   See "Configuring the Extended DHCP Tracing Flags" on page 258.

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 259.

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

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 276
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 11

DHCP Relay Agent Examples

- Example: Minimum DHCP Relay Agent Configuration on page 341
- Example: DHCP Relay Agent Configuration with Multiple Clients and Servers on page 342
- Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 343
- Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent on page 345
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 349
- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 353
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

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 10.0.2.1;
   }
   active-server-group test;
   group all {
      interface fe-0/0/2.0;
   }
}
```

NOTE: The interface type in this topic is just an example. The fe- interface type is not supported by EX Series switches.

This example creates a server group and an active server group named **test** with IP address 10.0.2.1. 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 276

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 {
      10.0.2.1;
      10.0.2.2;
    }
    sp-2 {
      10.33.2.1;
      10.33.2.2;
      10.33.2.3;
    }
  }
  active-server-group sp-1;
  overrides layer2-unicast-replies;
  group clients_a {
    relay-option-82 circuit-id;
    interface fe-1/0/1.1;
    interface fe-1/0/1.2;
    interface fe-1/0/1.3;
  }
  group clients_b {
    relay-option-82 {
      circuit-id {
        prefix routing-instance-name;
      }
    }
    interface fe-1/0/1.4;
    interface fe-1/0/1.5;
    interface fe-1/0/1.6;
  }
  group eth_dslam_relay {
    active-server-group sp-2;
    overrides {
      trust-option-82;
      layer2-unicast-replies;
    }
    interface fe-1/0/1.7;
    interface fe-1/0/1.8;
    interface fe-1/0/1.9;
  }
}
```

This example creates two server-groups: **sp-1**, which includes DHCP server addresses 10.0.2.1 and 10.0.2.2, and **sp-2**, which includes DHCP server addresses 10.33.2.1, 10.33.2.2, and 10.33.2.3. The active server group to which the DHCP relay agent configuration applies

 Copyright © 2013, Juniper Networks, Inc.

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 276

Example: Configuring DHCP Snooping Support for DHCP Relay Agent

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

- Requirements on page 343
- Overview on page 343
- Configuration on page 344

Requirements

- Configure DHCP relay agent. See "Extended DHCP Relay Agent Overview" on page 276.

Overview

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

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

NOTE: **By default, DHCP snooping is enabled globally in Junos OS Release 10.0 and earlier and disabled globally in Junos OS Release 10.1 and later.**

Configuration

Step-by-Step Procedure

To configure DHCP relay support for DHCP snooping:

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

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

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

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

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

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

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

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

5. Enable DHCP snooping support for the group.

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

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

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

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

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

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

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

Results

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

```
[edit]
regress@montag# show forwarding-options
dhcp-relay {
   forward-snooped-clients configured-interfaces;
   group frankfurt {
```


Copyright © 2013, Juniper Networks, Inc.

```
        overrides {
          allow-snooped-clients;
        }
        interface fe-1/0/1.3 {
          upto fe-1/0/1.9;
        }
      }
    }
```

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

Related Documentation

- DHCP Snooping Support on page 299
- Configuring DHCP Snooping for DHCP Relay Agent on page 301

Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent

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

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

- Requirements on page 345
- Overview on page 346
- Configuration on page 346
- Verification on page 348

Requirements

This example uses the following hardware and software components:

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

Before you begin:

- Configure DHCPv6 relay agent.

  See "DHCPv6 Relay Agent Overview" on page 279

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

  See "Grouping Interfaces with Common DHCP Configurations" on page 217

Overview

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

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

Configuration

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

- Enabling DHCPv6 Snooping Support for a Named Group of Interfaces on page 346
- Enabling DHCPv6 Snooping Support for a Specific Interface in a Named Group on page 347

CLI Quick Configuration

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

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

Enabling DHCPv6 Snooping Support for a Named Group of Interfaces

Step-by-Step Procedure

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

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

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

2. Specify the named group of interfaces for which you want to enable DHCPv6 snooping.

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

3. Specify that you want to override the default DHCPv6 configuration for the interfaces in that group.

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

4. Enable DHCPv6 snooping support for all interfaces in group **boston**.

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


Copyright © 2013, Juniper Networks, Inc.

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

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

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

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

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

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

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

2. Specify the named group containing the interface.

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

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

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

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

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

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

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

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

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

```
user@host# show
dhcpv6 {
  group boston {
    overrides {
      allow-snooped-clients;
    }
  }
  group sunnyvale {
    interface ge-3/2/1.1 {
      overrides {
        allow-snooped-clients;
      }
    }
  }
}
```

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

Verification

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

- Verifying the Address Bindings for DHCPv6 Relay Agent Clients on page 348

Verifying the Address Bindings for DHCPv6 Relay Agent Clients

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

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

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

Session Id:  13
     Client IPv6 Prefix:                    3000:0:0:8001::5/128
     Client DUID:                           LL0x1-00:00:65:03:01:02
     State:                                 BOUND(DHCPV6_RELAY_STATE_BOUND)
     Lease Expires:                         2011-11-21 06:14:50 PST
     Lease Expires in:                      293 seconds
     Lease Start:                           2011-11-21 06:09:50 PST
     Incoming Client Interface:             ge-3/2/1.1
     Server Address:                        unknown
     Next Hop Server Facing Relay:          4000::2
     Server Interface:                      none
     Client Id Length:                      10
     Client Id:                             /0x00030001/0x00006503/0x0102
```

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

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

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- DHCPv6 Relay Agent Overview on page 279
- DHCP Snooping Support on page 299
- Grouping Interfaces with Common DHCP Configurations on page 217
- Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 301

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 349
- Overview on page 349
- Configuration on page 350
- Verification on page 351

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Routers or EX Series Switches
- Junos OS Release 12.3 or later or Junos OS Release 12.3R2 for 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 276.

- (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 217.

Overview

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

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

Copyright © 2013, Juniper Networks, Inc.

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 350
- Results on page 351

CLI Quick Configuration

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

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

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

Step-by-Step Procedure

To configure DHCP relay selective processing:

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

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

2. Specify the DHCP option that DHCP relay agent uses to identify incoming client traffic.

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

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

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

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

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

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

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


Copyright © 2013, Juniper Networks, Inc.

```
user@host# set relay-option equals ascii video-bronze local-server-group
  servergroup-15
```

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

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

Results

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

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

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

Verification

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 351

Verifying the Status of DHCP Relay Agent Selective Traffic Processing

Purpose Verify the DHCP relay agent selective traffic processing status.

Action Display statistics for DHCP relay agent.

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

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

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

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

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

Related Documentation

- Extended DHCP Relay Agent Overview on page 276
- DHCP Options and Selective Traffic Processing Overview on page 321
- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 323
- Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings on page 325
- Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing on page 353

 Copyright © 2013, Juniper Networks, Inc.

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 353
- Overview on page 353
- Configuration on page 354
- Verification on page 356

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Routers or PTX Series Packet Transport Routers
- Junos OS Release 12.3 or later.

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

- Configure DHCPv6 relay agent.

  See "Extended DHCP Relay Agent Overview" on page 276 and "DHCPv6 Relay Agent Overview" on page 279.

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

  See "Grouping Interfaces with Common DHCP Configurations" on page 217.

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

  See "Grouping Interfaces with Common DHCP Configurations" on page 217.

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.

Copyright © 2013, Juniper Networks, Inc.

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 354
- Results on page 355

CLI Quick Configuration

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

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

Configuring a DHCPv6 Relay Agent Named Interface Group To Selectively Process Client Traffic Based on DHCPv6 Option Strings

Step-by-Step Procedure

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

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

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

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

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

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

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

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

 Copyright © 2013, Juniper Networks, Inc.

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

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

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

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

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

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

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

Results

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

```
dhcpv6 {
  group test-1 {
    relay-option {
      option-number 15;
      equals {
        ascii triple-gold {
          relay-server-group relayserver-triple-8;
        }
        ascii triple-silver {
          relay-server-group relayserver-triple-23;
        }
      }
      default-action {
        relay-server-group relayserver-def-4;
      }
      starts-with {
        ascii single {
          relay-server-group relayserver-1-aa;
        }
      }
    }
    interface ge-1/0/0.0 upto ge-1/1/0.0;
  }
  server-group {
    relayserver-1-aa;
    relayserver-triple-8;
```

```
        relayserver-triple-23;
        relayserver-def-4;
      }
    }
```

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

Verification

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

- Verifying the Status of DHCPv6 Relay Agent Selective Traffic Processing on page 356

Verifying the Status of DHCPv6 Relay Agent Selective Traffic Processing

Purpose Verify the DHCPv6 relay agent selective traffic processing status.

Action Display statistics for DHCPv6 relay agent.

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                        0

Messages received:
    DHCPV6_DECLINE               0
    DHCPV6_SOLICIT               10
    DHCPV6_INFORMATION_REQUEST   0
    DHCPV6_RELEASE               0
    DHCPV6_REQUEST               10
    DHCPV6_CONFIRM               0
    DHCPV6_RENEW                 0
    DHCPV6_REBIND                0
    DHCPV6_RELAY_REPL            0

Messages sent:
    DHCPV6_ADVERTISE             0
    DHCPV6_REPLY                 0
    DHCPV6_RECONFIGURE           0
    DHCPV6_RELAY_FORW            0

Packets forwarded:
    Total                        4
    FWD REQUEST                  2
    FWD REPLY                    2
```

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

Related Documentation

- Extended DHCP Relay Agent Overview on page 276
- DHCPv6 Relay Agent Overview on page 279
- DHCP Options and Selective Traffic Processing Overview on page 321

 Copyright © 2013, Juniper Networks, Inc.

- Using DHCP Option Information to Selectively Process DHCP Client Traffic on page 323
- Grouping Interfaces with Common DHCP Configurations on page 217
- Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings on page 325
- Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 349

Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

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

- Requirements on page 357
- Overview on page 357
- Configuration on page 358

Requirements

- Juniper Networks MX Series routers.
- Junos OS Release 12.1R1 or later.
- Configure DHCP relay agent. See "Extended DHCP Relay Agent Overview" on page 276.

Overview

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

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

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

Configuration

Step-by-Step Procedure

To configure liveness detection for DHCP relay:

1. Specify that you want to configure liveness detection.

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

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

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

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

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

4. Configure the detection time threshold (in milliseconds) at which a trap is produced.

   [edit forwarding-options dhcp-relay liveness-detection method bfd]
   user@host# **set** detection-time threshold **50000**

5. Configure the time (in milliseconds) for which BFD holds a session up notification.

   [edit forwarding-options dhcp-relay liveness-detection method bfd]
   user@host# **set** holddown-interval **50**

6. Configure the BFD minimum transmit and receive interval (in milliseconds).

   [edit forwarding-options dhcp-relay liveness-detection method bfd]
   user@host# **set** minimum-interval **45000**

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

   [edit forwarding-options dhcp-relay liveness-detection method bfd]
   user@host# **set** minimum-receive-interval **60000**

8. Configure a multiplier value for the detection time.

   [edit forwarding-options dhcp-relay liveness-detection method bfd]
   user@host# **set** multiplier **100**

9. Disable the ability for BFD interval timers to change or adapt to network situations.

   [edit forwarding-options dhcp-relay liveness-detection method bfd]
   user@host# **set** no-adaptation

10. Configure the BFD session mode.

    [edit forwarding-options dhcp-relay liveness-detection method bfd]
    user@host# **set** session-mode **automatic**

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

    [edit forwarding-options dhcp-relay liveness-detection method bfd]
    user@host# **set** transmit-interval threshold **60000** minimum-interval **45000**

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

    [edit forwarding-options dhcp-relay liveness-detection method bfd]
    user@host# **set** version **automatic**

Copyright © 2013, Juniper Networks, Inc.

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

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

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

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

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

Related Documentation

- Extended DHCP Relay Agent Overview on page 276
- DHCP Liveness Detection Overview on page 250
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 318

Copyright © 2013, Juniper Networks, Inc.

PART 5

PPP for Subscriber Access

- Dynamic Profiles for PPP Overview on page 363
- Configuring PPP for Subscriber Access on page 365
- Configuring Subscriber Services for MLPPP Interfaces on page 373

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 12

Dynamic Profiles for PPP Overview

- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 363

Dynamic Profiles for PPP Subscriber Interfaces Overview

Subscriber management PPP support enables you to create and attach dynamic profiles for PPP subscriber interfaces. When the PPP subscriber logs in, the router instantiates the specified dynamic profile and then applies the attributes defined in the profile to the interface.

Dynamic profiles are used for both static and dynamic PPP interfaces. For static PPP interfaces, you use the CLI to attach dynamic profiles, which specify PPP options. For dynamic PPP interfaces, the dynamic profile creates the interface, including the PPP options.

NOTE: **Dynamically created interfaces are supported only on PPPoE interfaces.**

Unlike traditional PPP support, subscriber management does not allow bi-directional PPP authentication—authentication is performed only by the router, never by the remote peer. The router's AAA process manages authentication and address assignment for subscriber management. When you configure PPP options for a dynamic profile, you can configure either Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP) authentication, and you can control the order in which the router negotiates the CHAP and PAP protocols. In addition, for CHAP authentication, you can modify the default length of the CHAP challenge message. Other PPP options, which are either commonly used or mandatory for a traditional PPP interface configuration, are not supported in subscriber management dynamic profiles.

Related Documentation

- Configuring Dynamic Authentication for PPP Subscribers on page 365
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 370
- Verifying and Managing PPP Configuration for Subscriber Management on page 372
- Example: Minimum PPPoE Dynamic Profile on page 697

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 13

Configuring PPP for Subscriber Access

- Configuring Dynamic Authentication for PPP Subscribers on page 365
- Controlling the Negotiation Order of PPP Authentication Protocols on page 367
- Modifying the CHAP Challenge Length on page 369
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 370
- Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests on page 371
- Verifying and Managing PPP Configuration for Subscriber Management on page 372

Configuring Dynamic Authentication for PPP Subscribers

You can configure a dynamic profile that includes PPP authentication that enables PPP clients to dynamically access the network. You can specify either CHAP or PAP authentication. Optionally, you can also control the order in which the router negotiates the CHAP and PAP protocols.

For dynamic interfaces, the router supports unidirectional authentication only—the router always functions as the authenticator. When you configure PPP authentication in a dynamic profile, CHAP authentication supports the **challenge-length** option, which enables you to configure the minimum length and maximum length of the CHAP challenge message. Neither CHAP authentication nor PAP authentication supports any other configuration options, including the **passive** statement.

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

To configure authentication in a dynamic profile for PPP subscriber interfaces:

1. Name the dynamic profile.

   ```
   [edit]
   user@host# edit dynamic-profiles vod-profile-25
   ```

2. Configure the interfaces and unit for the dynamic profile. Use **pp0** for the interface type and the Junos predefined variable for the unit.

   ```
   [edit dynamic-profiles vod-profile-25]
   ```

Copyright © 2013, Juniper Networks, Inc.

```
user@host# edit interfaces pp0 unit $junos-interface-unit
```

3. Configure PPP options.

```
[edit dynamic-profiles vod-profile-25 interfaces pp0 unit "$junos-interface-unit"]
user@host# edit ppp-options
```

4. Specify the authentication protocol used in the dynamic profile. You can configure either CHAP or PAP. There are no additional options for either authentication protocol.

```
[edit dynamic-profiles vod-profile-25 interfaces pp0 unit "$junos-interface-unit"
  ppp-options]
user@host# set chap
```

5. (Optional) Configure the minimum length and maximum length of the CHAP challenge message.

   See "Modifying the CHAP Challenge Length" on page 369.

6. (Optional) Configure the order in which the router negotiates the CHAP and PAP authentication protocols.

   See "Controlling the Negotiation Order of PPP Authentication Protocols" on page 367.

Related Documentation

- Modifying the CHAP Challenge Length on page 369
- Controlling the Negotiation Order of PPP Authentication Protocols on page 367
- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 363
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 370
- Dynamic Profiles Overview on page 638
- Configuring a Basic Dynamic Profile on page 673
- Example: Minimum PPPoE Dynamic Profile on page 697
- Verifying and Managing PPP Configuration for Subscriber Management on page 372

Copyright © 2013, Juniper Networks, Inc.

Controlling the Negotiation Order of PPP Authentication Protocols

You can control the order in which the router tries to negotiate PPP authentication protocols when it verifies that a PPP client can access the network. By default, the router first tries to negotiate Challenge Handshake Authentication Protocol (CHAP) authentication. If the the attempt to negotiate CHAP authentication is unsuccessful, the router then tries to negotiate Password Authentication Protocol (PAP) authentication.

You can modify this default negotiation order in any of the following ways:

- Specify that the router negotiate PAP authentication first, followed by CHAP authentication if PAP negotiation is unsuccessful.

  When you specify both authentication protocols in either order, you must enclose the set of protocol names in square brackets ([ ]).
- Specify that the router negotiate only CHAP authentication.
- Specify that the router negotiate only PAP authentication.

Before you begin:

- Configure the CHAP or PAP protocol on the interface.
  - For dynamic PPP subscriber interfaces, see "Configuring Dynamic Authentication for PPP Subscribers" on page 365.
  - For CHAP on static interfaces with PPP encapsulation, see *Configuring the PPP Challenge Handshake Authentication Protocol*.
  - For PAP on static interfaces with PPP encapsulation, see *Configuring the PPP Password Authentication Protocol*.

To control the order in which the router negotiates PPP authentication protocols:

1. Specify that you want to configure PPP options.
   - For dynamic PPP subscriber interfaces:

     [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit"]
     user@host# edit ppp-options
   - For static interfaces with PPP encapsulation:

     [edit interfaces pp0 unit *logical-unit-number*]
     user@host# edit ppp-options
2. Specify the negotiation order for PPP authentication protocols on the router.
   - For dynamic PPP subscriber interfaces:

     [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options]
     user@host# set authentication [*authentication-protocols*]
   - For static interfaces with PPP encapsulation:

     [edit interfaces pp0 unit *logical-unit-number* ppp-options]

```
user@host# set authentication [authentication-protocols]
```

The following sample **authentication** statements in a dynamic profile named pppoe-client-profile show the different ways you can configure the negotiation order for PPP authentication protocols. (The **authentication** statements for configuring static interfaces are identical.)

- To specify that the router negotiate PAP authentication first, followed by CHAP authentication:

  ```
  [edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit"
    ppp-options]
  user@host# set authentication [pap chap]
  ```

- To specify that the router negotiate only CHAP authentication:

  ```
  [edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit"
    ppp-options]
  user@host# set authentication chap
  ```

- To specify that the router negotiate only PAP authentication:

  ```
  [edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit"
    ppp-options]
  user@host# set authentication pap
  ```

- To restore the default negotiation order for PPP authentication protocols after you have modified it:

  ```
  [edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit"
    ppp-options]
  user@host# set authentication [chap pap]
  ```

Related Documentation

- Configuring Dynamic Authentication for PPP Subscribers on page 365
- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 363
- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Configuring the PPP Password Authentication Protocol*

Copyright © 2013, Juniper Networks, Inc.

Modifying the CHAP Challenge Length

You can modify the default minimum length and maximum length of the Challenge Handshake Authentication Protocol (CHAP) challenge message that the router sends to a PPP client. The CHAP challenge message, which contains information that is unique to a particular PPP subscriber session, is used as part of the authentication mechanism between the router and the client to verify the identity of the client for access to the router.

By default, the minimum length of the CHAP challenge is 16 bytes, and the maximum length is 32 bytes. You can override this default to configure the CHAP challenge minimum length and maximum length in the range 8 bytes through 63 bytes.

BEST PRACTICE: **We recommend that you configure both the minimum length and the maximum length of the CHAP challenge to at least 16 bytes.**

Before you begin:

- Configure the CHAP protocol on the interface.
  - For dynamic PPP subscriber interfaces, see "Configuring Dynamic Authentication for PPP Subscribers" on page 365.
  - For static interfaces with PPP encapsulation, see *Configuring the PPP Challenge Handshake Authentication Protocol*.

To configure the minimum and maximum length of the CHAP challenge message:

1. Specify that you want to configure PPP options.
   - For dynamic PPP subscriber interfaces:

     [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit"]
     user@host# edit ppp-options

   - For static interfaces with PPP encapsulation:

     [edit interfaces pp0 unit *logical-unit-number*]
     user@host# edit ppp-options

2. Specify that you want to configure CHAP options.
   - For dynamic PPP subscriber interfaces:

     [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options]
     user@host# edit chap

   - For static interfaces with PPP encapsulation:

     [edit interfaces pp0 unit *logical-unit-number* ppp-options]
     user@host# edit chap

3. Specify the minimum length and maximum length of the CHAP challenge.

- For dynamic PPP subscriber interfaces:

  [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options chap]
  user@host# set challenge-length minimum *minimum-length* maximum *maximum-length*

- For static interfaces with PPP encapsulation:

  [edit interfaces pp0 unit *logical-unit-number* ppp-options chap]
  user@host# set challenge-length minimum *minimum-length* maximum *maximum-length*

For example, the following **challenge-length** statement in a dynamic profile named pppoe-client-profile sets the minimum length of the CHAP challenge to 20 bytes, and the maximum length to 40 bytes.

```
[edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit"
  ppp-options chap]
user@host# set challenge-length minimum 20 maximum 40
```

Related Documentation

- Configuring Dynamic Authentication for PPP Subscribers on page 365
- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 363
- *Configuring the PPP Challenge Handshake Authentication Protocol*

Attaching Dynamic Profiles to Static PPP Subscriber Interfaces

You can attach a dynamic profile to a static PPP subscriber interface. When a PPP subscriber logs in, the specified dynamic profile is instantiated and the services defined in the profile are applied to the interface.

To attach a dynamic profile to a static PPP subscriber interface:

1. Specify that you want to configure PPP options.

   [edit interfaces pp0 unit 0]
   user@host# edit ppp-options

2. Specify the dynamic profile you want to associate with the interface.

   [edit interfaces pp0 unit 0 ppp-options]
   user@host# set dynamic-profile vod-profile-50

Related Documentation

- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 363
- Configuring Dynamic Authentication for PPP Subscribers on page 365
- Dynamic Profiles Overview on page 638
- Configuring a Basic Dynamic Profile on page 673
- Example: Minimum PPPoE Dynamic Profile on page 697
- Verifying and Managing PPP Configuration for Subscriber Management on page 372

Copyright © 2013, Juniper Networks, Inc.

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

On MX Series routers with Modular Port Concentrators/Modular Interface Cards (MPCs/MICs), the Packet Forwarding Engine on an MPC/MIC processes and responds to Link Control Protocol (LCP) Echo-Request packets that the PPP subscriber (client) initiates and sends to the router. LCP Echo-Request packets and LCP Echo-Reply packets are part of the PPP keepalive mechanism that helps determine whether a link is functioning properly.

Previously, LCP Echo-Request packets and LCP Echo-Reply packets were handled on an MX Series router by the Routing Engine. Support for the PPP fast keepalive feature enables the Packet Forwarding Engine on the MPC/MIC to receive LCP Echo-Request packets from the PPP subscriber and transmit LCP Echo-Reply packets in response, without having to send the LCP packets to the Routing Engine for processing. The mechanism by which LCP Echo-Request packets are processed by the Packet Forwarding Engine instead of by the Routing Engine is referred to as *PPP fast keepalive*.

Relieving the Routing Engine of having to process LCP Echo-Request packets provides increased bandwidth on the router to support a larger number of subscribers with improved performance.

- How PPP Fast Keepalive Processing Works on page 371
- Statistics Display for PPP Fast Keepalive on page 372
- Effect of Changing the Forwarding Class Configuration on page 372

How PPP Fast Keepalive Processing Works

You do not need any special configuration on an MX Series router with MPCs/MICs to enable processing of PPP fast keepalive requests on the Packet Forwarding Engine. The feature is enabled by default, and cannot be disabled.

The following sequence describes how an MX Series router processes LCP Echo-Request packets and LCP Echo-Reply packets on the Packet Forwarding Engine on the MPC/MIC:

1. The Routing Engine notifies the Packet Forwarding Engine when transmission of keepalive requests is enabled on a PPP logical interface. The notification includes the magic numbers of both the server and the remote client.
2. The Packet Forwarding Engine receives the LCP Echo-Request packet initiated by the PPP subscriber (client).
3. The Packet Forwarding Engine validates the peer magic number in the LCP Echo-Request packet, and transmits the corresponding LCP Echo-Reply packet containing the magic number negotiated by the router.
4. If the Packet Forwarding Engine detects a loop condition in the link, it sends the LCP Echo-Request packet to the Routing Engine for further processing.

   The Routing Engine continues to process LCP Echo-Request packets until the loop condition is cleared.

Copyright © 2013, Juniper Networks, Inc.

Transmission of keepalive requests from the Packet Forwarding Engine on the router is not currently enabled.

Statistics Display for PPP Fast Keepalive

When an MX Series router with MPCs/MICs is using PPP fast keepalive for a PPP link, the **Keepalive statistics** field in the output of the **show interfaces pp0.*logical* statistics** operational command does not include statistics for the number of keepalive packets received or sent, or the amount of time since the router received or sent the last keepalive packet.

Effect of Changing the Forwarding Class Configuration

To change the default queue assignment (forwarding class) for outbound traffic generated by the Routing Engine, you can include the **forwarding-class *class-name*** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level.

For PPP fast (inline) keepalive LCP Echo-Request and LCP Echo-Reply packets transmitted between an MX Series router with MPCs/MICs and a PPP client, changing the forwarding class configuration takes effect immediately for both new PPP-over-Ethernet (PPPoE), PPP-over-ATM (PPPoA), and L2TP network server (LNS) subscriber sessions created after the configuration change, and for existing PPPoE, PPPoA, and LNS subscriber sessions established before the configuration change.

Related Documentation

- *Configuring Keepalives*
- *Disabling the Sending of PPPoE Keepalive Messages* in *Ethernet Interfaces*
- *Changing the Default Queuing and Marking of Host Outbound Traffic* in *Junos OS Class of Service Library for Routing Devices*

Verifying and Managing PPP Configuration for Subscriber Management

Purpose View or clear information about PPP configuration for subscriber management.

Action

- To display information about PPP interfaces:

  user@host> **show ppp interface**

- To display PPP statistics information:

  user@host> **show ppp statistics**

- To display PPP session summary information:

  user@host> **show ppp summary**

Related Documentation

- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 363
- *Junos OS Operational Mode Commands*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 14

Configuring Subscriber Services for MLPPP Interfaces

- Dynamic PPP Subscriber Services for Static MLPPP Interfaces on page 373
- Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces on page 374
- Configuring PPP Subscriber Services for MLPPP Bundles on page 374
- Enabling PPP Subscriber Services for Static Non-Ethernet Interfaces on page 375
- Attaching Dynamic Profiles to MLPPP Bundles on page 375

Dynamic PPP Subscriber Services for Static MLPPP Interfaces

Dynamic subscriber services are supported for MLPPP bundle interfaces, with certain interface and hardware restrictions. See "Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces" on page 374. Multiclass MLPPP (MCML) enables the relative prioritization of up to eight classes of traffic over an MLPPP bundle, but only on link services intelligent queuing (IQ) (LSQ) interfaces.

RADIUS previously supported only authentication for MLPPP. Address management, service deactivation, and dynamic selection of subscriber properties based on RADIUS user ID are now also supported.

RADIUS can dynamically allocate IPv4 addresses for MLPPP connections. When the first subscriber logs in, an address is allocated. The same address is allocated to all links in a bundle. Any other address provided for any of the links is ignored. The IP address is released for re-allocation when the last member link in a bundle logs out. Similar to the address allocation, the services configured for the first subscriber to log in are configured for all subsequent subscribers in the bundle.

The Acct-Multi-Session-Id [50] attribute enables RADIUS to link multiple related sessions into a single log file. RADIUS uses the session database (SDB) bundle session ID for the value of Acct-Multi-Session-Id. This bundle ID enables RADIUS to initiate a disconnect for an entire bundle. By tracking the member link sessions, RADIUS is also able to disconnect the individual member links in a bundle.

The Acct-Link-Count [51] attribute records the number of links present in a multilink session at the time the accounting record is generated.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces on page 374
- Configuring PPP Subscriber Services for MLPPP Bundles on page 374

Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces

PPP subscriber services are supported for MLPPP bundle interfaces. These services require the following hardware:

- M120 router or M320 router
- Channelized DS3/E3 Enhanced IP PIC (PB-4CHDS3-E3-IQE-BNC) to support MLPPP subscriber access
- An Adaptive Services PIC or Multiservices PIC to support subscriber services on LSQ MLPPP bundle interfaces

Subscriber services are not supported for single-link PPP interfaces with this hardware.

Related Documentation

- Dynamic PPP Subscriber Services for Static MLPPP Interfaces on page 373

Configuring PPP Subscriber Services for MLPPP Bundles

You can configure PPP subscriber services for static LSQ MLPPP bundle interfaces.

To configure PPP subscriber services for static LSQ MLPPP bundle interfaces:

1. Enable PPP subscriber services for the interfaces.

   See "Enabling PPP Subscriber Services for Static Non-Ethernet Interfaces" on page 375.

2. Attach a dynamic profile to the MLPPP bundle interface.

   See "Attaching Dynamic Profiles to MLPPP Bundles" on page 375.

Related Documentation

- Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces on page 374
- Example: Minimum MLPPP Dynamic Profile on page 697
- Example: Configuring CoS on Static LSQ MLPPP Bundle Interfaces on page 811

Copyright © 2013, Juniper Networks, Inc.

Enabling PPP Subscriber Services for Static Non-Ethernet Interfaces

You can enable PPP subscriber services for certain non-Ethernet interface types on particular associated PICs. Supported interfaces are listed in "Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces" on page 374.

To enable PPP subscriber services on supported non-Ethernet interfaces:

- Configure PPP subscriber services.

```
[edit chassis]
user@host# set ppp-subscriber-services enable
```

To disable PPP subscriber services on supported non-Ethernet interfaces:

- Disable PPP subscriber services.

```
[edit chassis]
user@host# set ppp-subscriber-services disable
```

Related Documentation

- For hardware requirements, see Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces on page 374
- Configuring PPP Subscriber Services for MLPPP Bundles on page 374

Attaching Dynamic Profiles to MLPPP Bundles

You can attach a dynamic profile to a static MLPPP bundle interface. When a PPP subscriber logs in on a member link, the specified dynamic profile is instantiated and the services defined in the profile are applied to the LSQ bundle interface.

To attach a dynamic profile to a static LSQ MLPPP bundle interface:

1. Specify that you want to configure PPP options.

```
[edit interfaces lsq-3/3/0 unit 0]
user@host# edit ppp-options
```

2. Specify the dynamic profile you want to associate with the interface.

```
[edit interfaces lsq-3/3/0 unit 0 ppp-options]
user@host# set dynamic-profile vod-profile-50
```

Related Documentation

- Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces on page 374
- Configuring PPP Subscriber Services for MLPPP Bundles on page 374
- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 363
- Dynamic Profiles Overview on page 638
- Configuring a Basic Dynamic Profile on page 673
- Configuring PPP Subscriber Services for MLPPP Bundles on page 374

Copyright © 2013, Juniper Networks, Inc.

- Example: Minimum MLPPP Dynamic Profile on page 697
- Example: Configuring CoS on Static LSQ MLPPP Bundle Interfaces on page 811

Copyright © 2013, Juniper Networks, Inc.

PART 6

L2TP for Subscriber Access

- L2TP for Subscriber Access Overview on page 379
- Configuring L2TP for Subscriber Access on page 401

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 15

L2TP for Subscriber Access Overview

- L2TP for Subscriber Access Overview on page 379
- L2TP Terminology on page 382
- L2TP Implementation on page 383
- LAC Tunnel Selection Overview on page 384
- L2TP Failover and Peer Resynchronization on page 388
- L2TP and Graceful Routing Engine Switchover on page 389
- IP Packet Fragment Reassembly for L2TP Overview on page 390
- Transmission of Tx Connect-Speed and Rx Connect-Speed AVPs from LAC to LNS on page 391
- L2TP Tunnel Switching Overview on page 392
- Tunnel Switching Actions for L2TP AVPs at the Switching Boundary on page 396
- LAC Interoperation with Third-Party LNS Devices on page 398

L2TP for Subscriber Access Overview

The Layer 2 Tunneling Protocol (L2TP) is a client-server protocol that allows the Point-to-Point Protocol (PPP) to be tunneled across a network. L2TP encapsulates Layer 2 packets, such as PPP, for transmission across a network. An L2TP access concentrator (LAC), configured on an access device, receives packets from a remote client and forwards them to an L2TP network server (LNS) on a remote network. The LNS functions as the logical termination point of the PPP session tunneled by the LAC from the remote client. Figure 6 on page 380 shows a simple L2TP topology.

Copyright © 2013, Juniper Networks, Inc.

Figure 6: Typical L2TP Topology

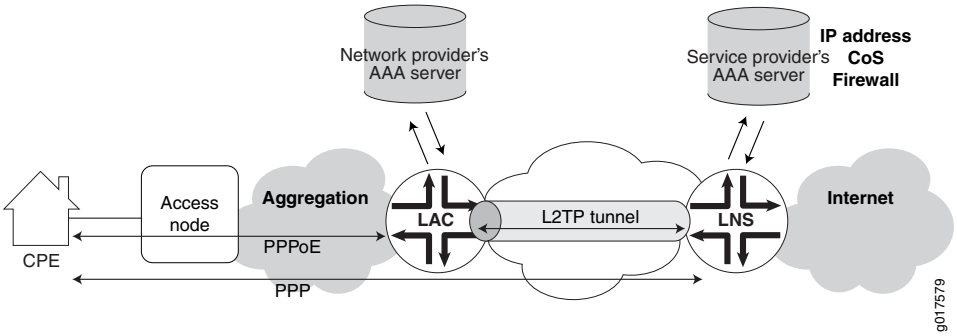

L2TP separates the termination of access technologies, such as cable or xDSL, from the termination of PPP and subsequent access to a network. This separation enables public ISPs to outsource their access technologies to competitive local exchange carriers (CLECs). L2TP provides ISPs the capability to supply VPN service; private enterprises can reduce or avoid investment in access technologies for remote workers.

You can configure your router to act as the LAC in PPP pass-through mode in which the LAC receives packets from a remote client and then forwards them at Layer 2 directly to the LNS. The PPP session is terminated on the LNS. This LAC implementation supports only Point-to-Point Protocol over Ethernet (PPPoE) subscribers over dynamic or static logical interfaces. Figure 7 on page 380 shows the protocol layer stacking for an L2TP pass-through connection.

Figure 7: Protocol Stacking for L2TP Subscribers in Pass-Through Mode

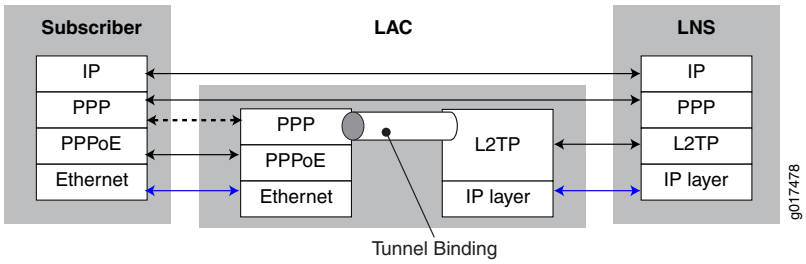

NOTE: On MX Series routers, the LAC and LNS functions are supported only on MPCs; they are not supported on any services PIC or MS-DPC. For details about MPC support for L2TP, see the *MX Series Interface Module Reference*

Certain M Series routers support LNS functions on services PICs. For more information about the L2TP implementation on M Series routers, see the *Junos OS Services Interfaces Library for Routing Devices*.

The LAC dynamically creates tunnels based on AAA authentication parameters and transmits L2TP packets to the LNS by means of the IP/User Datagram Protocol (UDP). Traffic travels in an L2TP *session*; a tunnel is an aggregation of one or more sessions. You can also provision a domain map that is used by AAA to determine whether to tunnel or

Copyright © 2013, Juniper Networks, Inc.

terminate the PPPoE subscriber on the LAC. A one-to-one mapping exists between each PPP subscriber tunneled to the LNS and an L2TP session.

When the LNS is an MX Series router, a LAC-facing peer interface on an MPC provides an IP address for the exchange of IP packets between the tunnel endpoints; the Routing Engine maintains the L2TP tunnels. The Packet Forwarding Engine hosts one or more inline services (**si**) interfaces. These interfaces function like a virtual physical interface and *anchor* the L2TP sessions on the LNS. The **si** interface enables L2TP services without requiring a special services PIC. Finally, another interface is used to transmit the subscriber data to and from the Internet.

The characteristics of the tunnel can originate either from a tunnel profile that you configure or from RADIUS tunnel attributes and vendor-specific attributes (VSAs) from the AAA server accessible at the LAC. You can include a tunnel profile in a domain map, which applies the tunnel profile before RADIUS authentication takes place. You can use RADIUS standard attributes and VSAs to override any or all characteristics configured by the tunnel profile in a domain map. Alternatively, RADIUS can itself apply a tunnel profile when the RADIUS Tunnel-Group VSA [26-64] is specified in the RADIUS login.

The Virtual-Router VSA [26-1] in the subscriber profile on the service provider AAA server (accessible from the LNS) determines the routing instance in which the L2TP session is brought up on the LNS. When this VSA is not present, the subscriber session comes up in the same routing instance as the tunnel, because the AAA server can be accessed only from the routing instance in which the tunnel terminates on the LNS.

This behavior is different than for DHCP and non-tunneled PPPoE subscribers, which come up in the default routing instance in the absence of the Virtual-Router VSA. For L2TP subscribers, you must include this VSA in the subscriber profile when you want the subscriber session to come up in a different routing instance than the tunnel routing instance.

The LAC supports RADIUS-initiated mirroring, which creates secure policies based on certain RADIUS VSAs, and uses RADIUS attributes to identify a subscriber whose traffic is to be mirrored. (This feature is not supported for an LNS configured on an MX Series router.)

The LAC supports unified ISSU. When an upgrade is initiated, the LAC completes any L2TP negotiations that are in progress but rejects any new negotiations until the upgrade has completed. No new tunnels or sessions are established during the upgrade. Subscriber logouts are recorded during the upgrade and are completed after the upgrade has completed. The LNS does not support unified ISSU. When LNS destinations exist, the LNS gracefully rejects the upgrade and the unified ISSU does not proceed.

Related Documentation

- RADIUS IETF Attributes Supported by the AAA Service Framework on page 91
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- Configuring a Tunnel Profile for Subscriber Access on page 403
- Domain Mapping Overview on page 180
- Subscriber Secure Policy and L2TP LAC Subscribers on page 1259

- *Unified ISSU Concepts*

L2TP Terminology

Table 40 on page 382 describes the basic terms for L2TP.

Table 40: L2TP Terms

Term	Description
AVP	Attribute value pair (AVP)—Combination of a unique attribute—represented by an integer—and a value containing the actual value identified by the attribute.
Call	A connection (or attempted connection) between a remote system and the LAC.
LAC	L2TP access concentrator (LAC)—A node that acts as one side of an L2TP tunnel endpoint and is a peer to the LNS. The LAC sits between an LNS and a remote system and forwards packets to and from each.
LNS	L2TP network server (LNS)—A node that acts as one side of an L2TP tunnel endpoint and is a peer to the LAC. The LNS is the logical termination point of a PPP connection that is being tunneled from the remote system by the LAC.
Peer	In the L2TP context, either the LAC or LNS. The LAC's peer is an LNS, and vice versa.
Proxy authentication	PPP pre-authentication performed by the LAC on behalf of the LNS. The proxy data is sent by the LAC to the LNS containing attributes such as authentication type, authentication name, and authentication challenge. The LNS responds with the authentication results.
Proxy LCP	Link Control Protocol (LCP) negotiation that is performed by the LAC on behalf of the LNS. The proxy is sent by the LAC to the LNS containing attributes such as the last configuration attributes sent and received from the client.
Remote system	An end system or router attached to a remote access network, which is either the initiator or recipient of a call.
Session	A logical connection created between the LAC and the LNS when an end-to-end PPP connection is established between a remote system and the LNS. NOTE: There is a one-to-one relationship between established L2TP sessions and their associated PPP connections.
Tunnel	A connection between the LAC-LNS pair consisting of a control connection and 0 or more L2TP sessions.

Related Documentation

- L2TP for Subscriber Access Overview on page 379

Copyright © 2013, Juniper Networks, Inc.

L2TP Implementation

L2TP is implemented on four levels:

- Source—The local router acting as the LAC.
- Destination—The remote router acting as the LNS.
- Tunnel—A direct path between the LAC and the LNS.
- Session—A PPP connection in a tunnel.

When the router has established destinations, tunnels, and sessions, you can control the L2TP traffic. Making a change to a destination affects all tunnels and sessions to that destination; making a change to a tunnel affects all sessions in that tunnel. For example, closing a destination closes all tunnels and sessions to that destination.

Sequence of Events on the LAC

The router acting as the LAC creates destinations, tunnels, and sessions dynamically, as follows:

1. The client initiates a PPP connection with the router.
2. The router and the client exchange Link Control Protocol (LCP) packets. The LAC negotiates on behalf of the LNS; this is known as *proxy LCP*.
3. The LAC authenticates the client on behalf of the LNS; this is known as *proxy authentication*. By using either a local database related to the domain name or RADIUS authentication, the router determines either to terminate or to tunnel the PPP connection.
4. If the router discovers that it should tunnel the session, it does the following:
   a. Sets up a new destination or selects an existing destination.
   b. Sets up a new tunnel or selects an existing tunnel.

      When a shared secret is configured in either the tunnel profile or the RADIUS attribute Tunnel-Password [69]—depending on which method is used to configure the tunnel—the secret is used to authenticate the tunnel during the establishment phase. The LAC includes the Challenge AVP in the SCCRQ message sent to the LNS. The LNS returns the Challenge Response AVP in the SCCRP message. If the response from the LNS does not match the value expected by the LAC, then tunnel authentication fails and the tunnel is not established.
   c. Opens a new session.
5. The router forwards the results of the LCP negotiations and authentication to the LNS.

A PPP connection now exists between the client and the LNS.

Copyright © 2013, Juniper Networks, Inc.

NOTE: The router discards received packets if the size of the variable-length, optional offset pad field in the L2TP header is too large. The router always supports packets that have an offset pad field of up to 16 bytes, and may support larger offset pad fields, depending on other information in the header. This restriction is a possible, although unlikely, cause of excessive discarding of L2TP packets.

Sequence of Events on the LNS

A router acting as an LNS might be set up as follows:

1. The LAC initiates a tunnel with the router acting as the LNS.
2. The LNS verifies that a tunnel with this LAC is valid: the destination is configured, the hostname and the tunnel password are correct.
3. The LNS completes the tunnel setup with the LAC.
4. The LAC sets up a session and initiates a session request to the LNS.
5. The LNS uses a static interface or creates a dynamic interface to anchor the PPP session.
6. If they are enabled and present, the LNS accepts the proxy LCP and the proxy authentication data and passes them to PPP.
7. PPP processes the proxy LCP, if it is present, and, if the proxy LCP is acceptable, places LCP on the LNS in opened state without renegotiation of LCP.
8. PPP processes the proxy authentication data, if it is present, and passes the data to AAA for verification. (If the data is not present, PPP requests the data from the peer.)

   NOTE: When the proxy LCP is not present or not acceptable, the LNS negotiates LCP with the peer. When LCP renegotiation is enabled on the LNS, the LNS ignores any pre-negotiated LCP parameters and renegotiates both the LCP parameters and PPP authentication with the PPP client.

9. The LNS passes the authentication results to the peer.

Related Documentation

- L2TP for Subscriber Access Overview on page 379

LAC Tunnel Selection Overview

L2TP enables you to specify:

- Up to 31 destinations for a domain.
- Up to eight levels of preference. Preference indicates the order in which the router attempts to connect to the destinations specified for a domain. Zero (0) is the highest level of preference.

Copyright © 2013, Juniper Networks, Inc.

- Up to 31 destinations for a single preference level.

When the LAC determines that a PPP session should be tunneled, it selects a tunnel from a set of tunnels associated with either the PPP user or the PPP user's domain. The router provides the following methods for selecting tunnels:

- Tunnel selection failover between preference levels (the default behavior)
- Tunnel selection failover within a preference level
- Maximum sessions per tunnel
- Weighted load balancing

Tunnel Selection Failover Between Preference Levels

When a user tries to log in to a domain, in the default method, the router attempts to connect to a destination in that domain by means of the associated tunnel with the highest preference level. If more than one destination is considered reachable by a tunnel in the preference level, the router randomly selects a destination and attempts to contact it through its associated tunnel at that level. If the router is unsuccessful, it marks the destination as unreachable and does not try to connect to that destination for a period of time called the *destination lockout timeout*. The default timeout is 5 minutes. The router then moves to the next lower preference level and repeats the process.

For example, suppose that there are three destinations for a domain and a tunnel has been defined for each destination: A, B, and C. All destinations are considered reachable, and the preference levels for the tunnels are assigned as follows:

- A at preference 0
- B at preference 1
- C at preference 2

When a PPP user tries to connect to the domain, the router acts as follows:

1. The router initially attempts to reach a destination by a tunnel at preference level 0, destination A in this example.
2. If this connection attempt fails, the router excludes destination A from consideration for the length of the destination lockout timeout.
3. The router goes to the next level, preference level 1, to reach a destination for the domain. At preference level 1, it attempts to connect to destination B.
4. If the connection attempt to destination B fails, the router excludes destination B from consideration for the length of the destination lockout timeout.
5. The router goes to the next level, preference level 2, and attempts to connect to destination C, the only destination in the domain that is still available.

Copyright © 2013, Juniper Networks, Inc.

6. If that attempt also fails, the router excludes destination C from consideration for the length of the destination lockout timeout.
7. At this point, the router has attempted to connect to every destination through every tunnel available for the domain and cycles back to preference level 0. When the destination lockout timeout for destination A expires, the router can attempt again to connect to destination A, and so on.

Although the destination lockout timeout typically prevents an unreachable destination from being tried until the period expires, the timeout is ignored in some circumstances. For example, If all destinations at a preference level are marked as unreachable when a user tries to log in to a domain, the router chooses and attempts to connect to the destination that failed first and therefore has the shortest time remaining until the destination lockout timeout expires. The key is to understand that the router always chooses a single destination at each level of preference, even if all destinations have recently failed.

If more than one destination for the domain is present at a preference level, the router randomly selects among them. If the router fails to connect to a destination at all preference levels with destinations for the domain, it cycles back to the highest level that still has a destination not excluded by an attempt.

For example, suppose that again there are three destinations for a domain and a tunnel has been defined for each destination: A, B, and C. All destinations are considered reachable, but the tunnels are distributed among the preference levels as follows:

- A and B at preference 0
- C at preference 1

In this example, when a PPP user tries to connect to the domain, the router randomly selects between A and B at preference level 0 and chooses destination B.

1. If the connection attempt to destination B fails, the router excludes destination B from consideration for the length of the destination lockout timeout.
2. The router goes to the next level, preference level 1, to reach a destination for the domain. At preference level 1, it attempts to connect to destination C.
3. If the connection attempt to destination C also fails, the router excludes destination C from consideration for the length of the destination lockout timeout.
4. The router cycles back to preference level 0.
   - If destination B is still locked out, the router attempts to connect to destination A.
   - If the lockout for destination B has expired, then the selection process starts over and the router again randomly selects between A and B to attempt a connection.

Tunnel Selection Failover Within a Preference Level

When tunnel selection failover within a preference level is configured, if the router tries to connect to a destination and is unsuccessful, it selects a new destination at the same preference level. If all destinations at a preference level are marked as unreachable, the

 Copyright © 2013, Juniper Networks, Inc.

router does not attempt to connect to a destination at that level. It drops to the next lower preference level to select a destination.

If all destinations at all preference levels are marked as unreachable, the router chooses the destination that failed first and tries to make a connection. If the connection fails, the router rejects the PPP user session without attempting to contact the remote router.

For example, suppose that there are four destinations for a domain and a tunnel has been defined for each destination: A, B, C, and D. All destinations are considered reachable, and the preference levels for the tunnels are assigned as follows:

- A and B at preference 0
- C and D at preference 1

In this example, when a PPP user tries to connect to the domain, the router randomly selects between A and B at preference level 0 and chooses destination B.

1. If the connection attempt to destination B fails, the router excludes tunnel B from consideration for the length of the destination lockout timeout.
2. The router attempts to connect to destination A at preference level 0.
3. If the connection attempt to destination A fails, the router excludes destination A from consideration for the length of the destination lockout timeout.
4. The router goes to the next level, preference level 1, to reach a destination for the domain. At preference level 1, the router randomly selects tunnel C.
5. If the connection attempt to destination C fails, the router excludes destination C from consideration for the length of the destination lockout timeout.
6. The router attempts to connect to destination D at preference level 1.
7. If the connection attempt to destination D fails, the router attempts again to connect to destination B again, its original selection.
8. If this second connection attempt to destination B fails, the user session is rejected.

Tunnel Selection and Maximum Sessions per Tunnel

When the maximum number of sessions allowed per tunnel is configured, the router takes that setting into consideration during the tunnel selection process. The maximum number of sessions per tunnel can be configured through a RADIUS Tunnel-Max-Sessions VSA [26-64] or by including the **max-sessions** statement in a tunnel profile.

If a randomly selected tunnel has a current session count equal to its maximum session count, the router does not attempt to connect to a destination with that tunnel. Instead, it selects an alternate tunnel from the set of reachable tunnels at the same preference level. If no additional reachable tunnels exist at the current preference level, the router drops to the next lower preference level to make the selection. This process is consistent, regardless of which failover scheme is currently running on the router.

Copyright © 2013, Juniper Networks, Inc.

If the maximum number of sessions is not configured for a tunnel, then that tunnel has no upper limit on the number of sessions it can support. By default, the maximum sessions value is 0 (zero), which allows unlimited sessions in the tunnel.

Tunnel Selection with Weighted Load Balancing

The maximum sessions value for a tunnel is used for weighted load balancing to select among multiple tunnels with the same preference level.

The weight of a tunnel is determined by the tunnel's maximum session limit and the maximum session limits of the other tunnels at the same preference level. The tunnel with the largest maximum session value has the largest weight. The tunnel with the next largest maximum session value has the next largest weight, and so on. The tunnel with the smallest maximum session value has the smallest weight.

Related Documentation

- Configuring the L2TP LAC Tunnel Selection Parameters on page 406
- Configuring the L2TP Destination Lockout Timeout on page 414

L2TP Failover and Peer Resynchronization

L2TP failover enables a failed L2TP endpoint to resynchronize with its nonfailed peer during recovery and restart of the L2TP protocol on the failed endpoint. L2TP failover is enabled by default.

The failover and L2TP peer resynchronization process does all of the following:

- Prevents the nonfailed endpoint from prematurely terminating a tunnel while the failed endpoint is recovering.
- Reestablishes the sequence numbers required for the operation of the L2TP control protocol.
- Resolves inconsistencies in the tunnel and session databases of the failed endpoint and the nonfailed endpoint.

The router supports both the L2TP failover protocol method (described in *RFC 4951, Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*) and the L2TP silent failover method. The differences between these two methods are as follows:

- With the L2TP failover protocol method, both endpoints must support the method or recovery always fails. The L2TP failover protocol method also requires a nonfailed endpoint to wait an additional recovery time period while the failed endpoint is recovering to prevent the nonfailed endpoint from prematurely disconnecting the tunnel. The additional recovery period delays the detection of tunnel keepalive failures.
- Silent failover operates entirely within the failed endpoint and does not require nonfailed endpoint support—this improves interoperability between peers. Silent failover does not require additional recovery time by the nonfailed endpoint, which also eliminates the potential for degraded responsiveness to the loss of tunnel connectivity.

 Copyright © 2013, Juniper Networks, Inc.

The default resynchronization method in Junos OS is *failover-protocol-fall-back-to-silent-failover*. The recovery method used depends on the results of the failover capability negotiation that takes place between L2TP peers when they establish a tunnel, which works as follows:

- L2TP on the LAC by default attempts to negotiate the L2TP failover protocol first. When L2TP determines that the remote peer supports the L2TP failover protocol, then the L2TP failover protocol method is used.
- When L2TP determines that the remote peer does not support the L2TP failover protocol, then the L2TP silent failover method is used. Falling back on this secondary method prevents the failover from forcing a disconnection of the tunnel to the peer and all its sessions.

You can change the default behavior by including the **disable-failover-protocol** statement at the **[edit services l2tp]** hierarchy level. This statement forces the LAC to operate only in silent failover mode. This configuration can be useful when routers that act as the LNS either are configured for silent failover or incorrectly negotiate use of the failover protocol even though they do not support it. However, when you issue this statement and the LNS supports only failover protocol, then the LAC cannot negotiate failover protocol, and recovery (failover protocol recovery initiated by the LNS) always fails.

Related Documentation

- L2TP and Graceful Routing Engine Switchover on page 389
- L2TP for Subscriber Access Overview on page 379

L2TP and Graceful Routing Engine Switchover

Graceful Routing Engine switchover (GRES) is supported on MX Series routers acting as either the L2TP LAC or LNS. In the event that L2TP (the l2tp-universal-edge process) restarts or that the router fails over from the active routing engine (RE) to the standby RE, L2TP graceful Routing Engine switchover ensures that the following occurs:

- The LAC and the LNS recover destinations, tunnels, and sessions that were already established at the time of the failure or restart.
- The LAC and the LNS respond to tunnel keepalive requests received during the switchover for established tunnels, but do not generate any keepalives until the switchover is complete.
- The LAC and the LNS delete all the tunnels and sessions that are not in the Established state.
- The LAC and the LNS reject requests to create new tunnels and sessions.
- The LAC and the LNS send another disconnect notification to the peer for sessions and tunnels that are already in the Disconnecting state at the time of the failure or restart. For sessions and tunnels that were coming up at that time, the LAC and LNS send a disconnect notification to the peer.
- The LAC and the LNS restart timers for the full timeout period for recovered L2TP destinations, tunnels, and sessions.

Copyright © 2013, Juniper Networks, Inc.

NOTE: Graceful Routing Engine switchover is supported only by L2TP LAC and LNS on MX Series routers. It is not supported by L2TP LNS on M Series routers.

Related Documentation

- L2TP Failover and Peer Resynchronization on page 388
- L2TP for Subscriber Access Overview on page 379

IP Packet Fragment Reassembly for L2TP Overview

You can configure the service interfaces on the MX Series routers with modular port concentrators (MPCs) to support reassemble fragmented IP packets for an L2TP connection. When packets are transmitted over an L2TP connection, the packets may be fragmented during transmission and need to be reassembled before they are processed further. Efficient reassembly is important for network throughput, scalability, and graceful response to congestion.

Fragmentation of IP packets for transmission and the need to reassemble the IP packets at a destination is a feature of how Layer 2 (the frame layer) and Layer 3 (the packet layer) operate. The maximum size of a frame, set by the Maximum Transmission Unit (MTU) value, and the maximum size of a packet are determined independently. Typically the packet size can far exceed the MTU size defined for the outgoing connection. If the packet size (data plus IP and other headers) exceeds the configured frame size (usually set by the transport medium limits), the packet must be fragmented and split across multiple frames for transmission. Frames are always processed immediately, when they arrive (if error-free), but packet fragments cannot be processed until the whole packet has been reassembled. Each packet fragment inside a frame series, except the last packet fragment, has the more fragments (MF) IP header bit set, indicating that this packet is part of a whole. The last packet fragment inside a frame does not have this MF bit set and therefore ends the fragment sequence. After all of the fragments of a packet have arrived, the entire packet can be reassembled.

In an L2TP connection, packets are transmitted between the L2TP Access Concentrator (LAC) and the L2TP Network Server (LNS). For an IP packet being transmitted over an L2TP connection, the packet is fragmented at the LAC, at an LNS, or at any intermediate router. IP reassembly parameters configured on the service interfaces of the LAC and the LNS determine how the fragments are reassembled at the service interfaces to ensure efficient reassembly over an L2TP connection.

Related Documentation

- Configuring IP Inline Reassembly for L2TP on page 421
- *Protocols and Applications Supported by MX240, MX480, MX960, and MX2020 Enhanced MPCs (MPCEs)*
- ip-reassembly on page 1742

 Copyright © 2013, Juniper Networks, Inc.

Transmission of Tx Connect-Speed and Rx Connect-Speed AVPs from LAC to LNS

An L2TP access concentrator (LAC) uses Incoming-Call-Connected (ICCN) messages to send the attribute-value pair (AVP) 24 (which represents the transmit connect speed) and the AVP 38 (which represents the receive connect speed) to the L2TP network server (LNS).

The value of speed in AVP 24 and AVP 38 is typically not greater than the value that is enforced by CoS on the LAC side of the network. Any difference between the speed reported in these AVPs and that enforced by CoS can occur because of differences between the CoS configuration (of the source that is used to enforce a downstream speed) and the transmit connect speed method used to establish these AVPs.

When you include the **tx-connect-speed-method** statement at the **[edit services l2tp]** hierarchy level, the transmit connect speed method selected for the downstream speed, AVP 24, also applies to the selection of the upstream speed , AVP 38. You can configure the transmit and receive connect speeds to be derived from Point-to-Point Protocol over Ethernet (PPPoE) intermediate agent tags that are sent from the digital subscriber line access multiplexer (DSLAM) to the LAC, the Access Node Control Protocol (ANCP) settings of the underlying interface, or the recommended (advisory) shaping rate. When the method you specify as ANCP, PPPoE intermediate agent tags (PPPoE IA tags), or advisory shaping rate does not support an upstream speed, the static or advisory speed is used.

A fallback method is adopted to derive the transmit and receive connect speed values when these values cannot be determined from any of the configured methods, such as from the ANCP settings, the PPPoE IA tags, or the advisory shaping rate, or if the speed value is determined to be 0 from any of the configured methods.

If you configure the ANCP method to calculate the connect speed, the following sequence of events takes place:

1. The upstream and downstream connect speed values are derived from ANCP.
2. If the values cannot be derived from ANCP, the PPPoE IA tags are used to determine the values. If the PPPoE IA tags are present for either or both transmit and receive connect speeds, these values are used.
3. If the values cannot be derived from the PPPoE IA tags, the recommended (advisory) shaping rate configured on the PPPoE logical interface is used. If the advisory shaping rate is present for either or both transmit and receive connect speeds, these values are used.
4. If the values cannot be derived from the advisory shaping rate, the configured or default port speed is used for transmit and receive connect speeds.

If you configure the PPPoE IA tags method to calculate the connect speed, the following sequence of events takes place:

1. The upstream and downstream connect speed values are derived from PPPoE IA tags.

2. If the values cannot be derived from the PPPoE IA tags, the recommended (advisory) shaping rate configured on the PPPoE logical interface is used. If the advisory shaping rate is present for either or both transmit and receive connect speeds, these values are used.
3. If the values cannot be derived from the advisory shaping rate, the configured or default port speed is used for transmit and receive connect speeds.

If you configure the static or advisory downstream shaping rate method to calculate the connect speed, the following sequence of events takes place:

1. The upstream and downstream connect speed values are derived from the advisory shaping rate.
2. If the values cannot be derived from the advisory shaping rate, the default port speed is used for transmit and receive connect speeds.

The transmit connect speed, AVP 24, is set in the ICCN messages on the basis of the method for determining the transmit connect speed configured using the **tx-connect-speed-method** statement at the **[edit services l2tp]** hierarchy level. You can configure the method for determining the transmit connect speed in the following order of precedence:

1. **ancp**—The speed is derived from the configured ANCP value for the underlying interface. You can change this speed after a subscriber has logged in, but those changes do not affect the actual speed used by the LNS.
2. **pppoe-ia-tags**—PPPoE IA tags sent from the DSLAM to the LAC. This speed value is transmitted when a subscriber logs in and it cannot be subsequently changed. This value is used when the **ancp** value is not available. This speed does not apply to the subscribers that are already logged in; it is effective only for subscribers that log in after this setting has been saved.
3. **static**—The speed is derived from the recommended (advisory) shaping rate configured on the PPPoE logical interface underlying the subscriber interface. If the advisory shaping rate is not configured on the underlying interface, then the actual port speed is used.

If you do not configure the transmit connect speed using the CLI interface, and if the advisory speed is also not available, then the actual port speed is used. For ge and xe interfaces, the speed value is set to 10,000,000 and for ae interfaces, the speed value is set to 0 and sent in both AVP 24 and AVP 38.

Related Documentation

- Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal on page 409
- Configuring an L2TP LAC on page 402

L2TP Tunnel Switching Overview

L2TP tunnel switching, also known as L2TP multihop, simplifies the deployment of an L2TP network across multiple domains. A router that lies between a LAC and an LNS is

 Copyright © 2013, Juniper Networks, Inc.

configured as an *L2TP tunnel switch* (LTS)—sometimes referred to simply as a *tunnel switch* or a *tunnel switching aggregator* (TSA)—as shown in Figure 8 on page 393. The LTS is configured as both an LNS and a LAC. When a remote LAC sends encapsulated PPP packets to the LNS configured on the LTS, the LTS can forward or redirect the packets through a different tunnel to a different LNS beyond the LTS. The logical termination point of the original L2TP session is switched to a different endpoint.

For example, in the network shown in Figure 8 on page 393, packets from the subscriber provisioned by service provider A are initially targeted at the LNS configured on the LTS. The LTS might redirect those packets to LNS1.

Figure 8: L2TP Tunnel Switching Network Topology

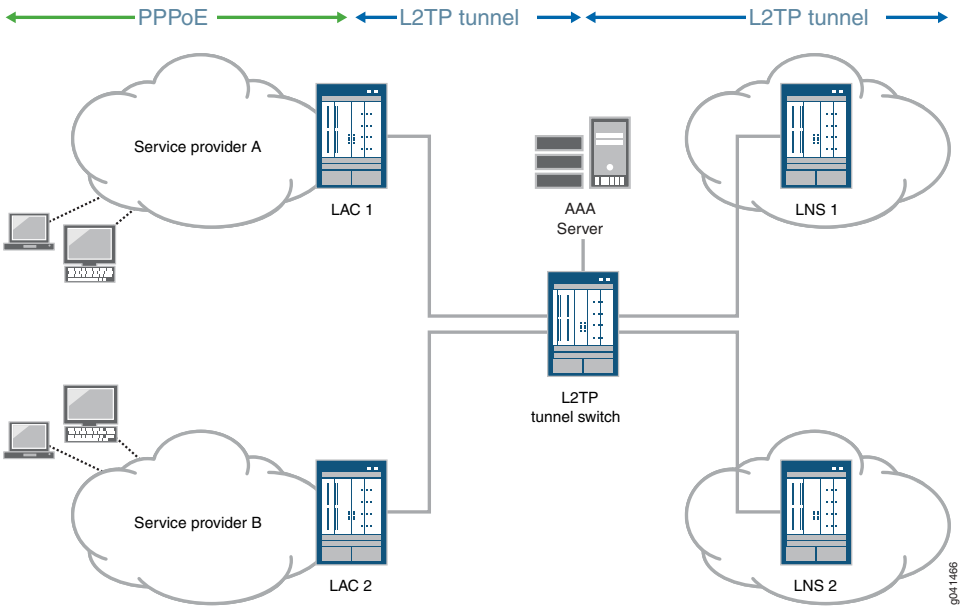

L2TP tunnel switching simplifies network configuration when the administrative domain of a LAC is different from that of the desired LNS. For example:

- The LTS acts as the LNS for multiple LACs. The individual LACs do not have to have the administrative control or capability required to identify the most appropriate LNS on which to terminate their sessions. The LTS performs that function is centralized in the LTS.
- The LTS acts as the LAC for multiple LNSs. When a new remote LAC is added to an ISP's network, the ISP does not have to reconfigure its LNS routers to accommodate the new LAC, because they connect to the LAC on the LTS.

In a Layer 2 wholesale network, the wholesaler can use L2TP tunnel switching to create a flatter network configuration that is easier to manage. The wholesaler bundles Layer 2 sessions from a LAC that are destined for different ISPs—and therefore different LNSs—onto a single L2TP tunnel. This configuration enables a common L2TP control connection to be used for the LAC.

Copyright © 2013, Juniper Networks, Inc.

Figure 9 on page 394 shows an example of L2TP tunnel switching for incoming calls with the following sequence of events:

1. The subscriber opens a PPP session to the LAC.
2. The LAC creates the first L2TP tunnel to the LNS configured on the LTS and the first L2TP session to carry the encapsulated PPP packets.
3. During authentication of this first session, the LTS determines whether to retunnel the session to an LNS beyond the LTS, based on the presence or absence of a tunnel switch profile configured on the LTS.

   The tunnel switch profile can be a default profile or it can be applied by the RADIUS server, a domain map configuration, or a tunnel group configuration.
4. If a tunnel switch profile is configured, the LTS creates a second tunnel (if it does not already exist) to the LNS beyond the LTS as specified in the profile and creates the second session in this tunnel.

Figure 9: L2TP Tunnel Switching for Incoming Calls

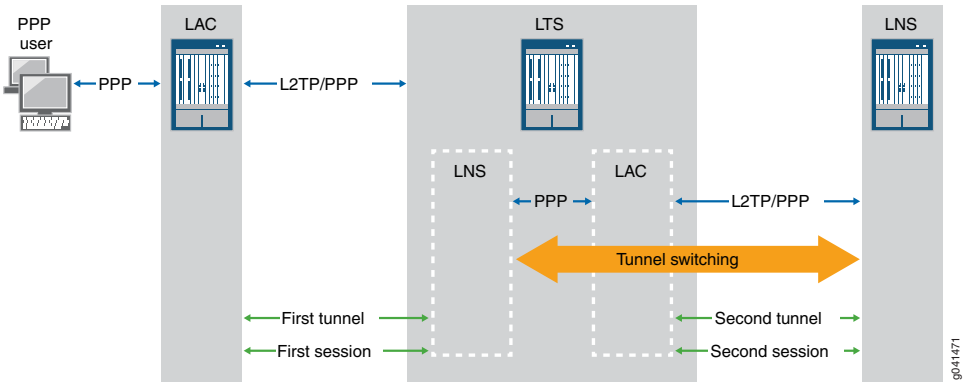

Application of Tunnel Switch Profiles

You can configure a tunnel switch profile to be applied in several ways:

- As a default profile applied globally to traffic received from all LACs
- With a domain map applied to a subscriber session
- With a tunnel group applied to a subscriber session
- In your RADIUS server configuration

You can configure more than one of these methods of application. When multiple tunnel switch profiles are present, the following order of precedence establishes which profile the LTS uses; the order is from highest (RADIUS) to lowest (default profile):

RADIUS > domain map > tunnel group > default tunnel switch profile

The tunnel switch profile must also reference a tunnel profile. This tunnel profile specifies the characteristics of the second tunnel, to which the subscriber packets are switched.

Copyright © 2013, Juniper Networks, Inc.

Termination of Tunnel-Switched Sessions on the LTS

Tunnel switched sessions are terminated on the LTS when any of the following happens:

- Either the LAC or LNS interface on the LTS receives a Call-Disconnect-Notify (CDN) message (Table 41 on page 395).

Table 41: Cause of CDN Message

CDN Message Is Received On	When
LAC interface	Either of the following occurs: • The second session cannot be established. • The remote LNS terminates the second session.
LNS interface	Either of the following occurs: • The PPPoE client initiates a logout. • The originating LAC initiates termination of the tunnel

Both the first and second sessions are terminated because the LTS relays the CDN to the interface that did not receive the CDN. The disconnect cause is the same for both sessions.

- Either the LAC or LNS interface on the LTS receives a Stop-Control-Connection-Notification (StopCCN) message (Table 42 on page 395).

Table 42: Cause of StopCCN Message

StopCCN Message Is Received On	When
LAC interface	Either of the following occurs: • The second session cannot be established. • The remote LNS terminates the second tunnel.
LNS interface	The originating LAC initiates termination of the tunnel.

The LTS does not relay the StopCCN message, because a given tunnel can contain both switched and nonswitched sessions. Another reason in a wholesale scenario is that the tunnel ending on the LNS on the LTS can contain sessions from LACs from different providers. Instead, the LTS sends a CDN message to the interface that did not receive the StopCCN to terminate the tunnel-switched session. This CDN relays the error code carried in the StopCCN.

- An administrative **clear** command is issued on the LTS.

Table 43 on page 396 lists the actions taken when an administrative **clear** command is issued on the LTS.

Copyright © 2013, Juniper Networks, Inc.

Table 43: LAC, LNS, and LTS Actions Taken for Switched Tunnels in Response to Administrative clear Commands

Command	LAC or LNS Action	LTS Action
clear services l2tp destination	Clear the destination and all associated tunnels and sessions.	For each switched session in a tunnel to the destination, clear the corresponding mapped switched session by sending it a CDN message with the cause set to Administrative.
clear services l2tp destination all	Clear all destinations and all associated tunnels and sessions.	None.
clear services l2tp session	Clear the session.	Clear the corresponding mapped switched session for this session by sending it a CDN message with the cause set to Administrative.
clear services l2tp session all	Clear all sessions.	None.
clear services l2tp tunnel	Clear the tunnel and all its sessions.	For each switched session in the tunnel, clear the corresponding mapped switched session by sending it a CDN message with the cause set to Administrative.
clear services l2tp tunnel all	Clear all tunnels.	None.

Related Documentation

- Configuring L2TP Tunnel Switching on page 432
- Tunnel Switching Actions for L2TP AVPs at the Switching Boundary on page 396
- L2TP for Subscriber Access Overview on page 379

Tunnel Switching Actions for L2TP AVPs at the Switching Boundary

When L2TP tunnel switching redirects packets to a different LNS, it performs one of the following default actions at the switching boundary for each AVP carried in the L2TP messages:

- **regenerate**—L2TP regenerates the AVP based on the local policy at the LTS and sent in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- **relay**—L2TP transparently forwards the AVP in the switched packet with no alteration.

Table 44 on page 397 lists the default action for each AVP. Mandatory AVPs are always included in the L2TP messages from the LAC; optional AVPs might be included in the messages.

You can optionally override the default action taken at the switching boundary for the Bearer Type AVP (18), Calling Number AVP (22), or Cisco NAS Port Info AVP (100). You can configure any of these three AVPs to be dropped from the switched packets or regenerated, or you can restore the default relay action.

 Copyright © 2013, Juniper Networks, Inc.

NOTE: L2TP AVPs that have their attribute values hidden are always regenerated at the switching boundary. The value is decoded and sent in clear text when the packet is forwarded to the remote LNS.

Table 44: Default Action for Handling L2TP AVPs at the Switching Boundary

AVP Name	AVP Type	L2TP Message Type	Default Action
Assigned Session Id	Mandatory	ICRQ, CDN	Regenerate
Assigned Tunnel Id	Mandatory	SCCRQ, CDN	Regenerate
Bearer Capabilities	Optional	SCCRQ	Regenerate
Bearer Type	Optional	ICRQ	Relay
Call Serial Number	Mandatory	ICRQ	Relay
Called Number	Optional	ICRQ	Relay
Calling Number	Optional	ICRQ	Relay
Challenge	Optional	SCCRQ	Regenerate
Challenge Response	Optional	SCCCN	Regenerate
Cisco NAS Port	Optional	ICRQ	Relay
Failover Capability	Optional	SCCRQ	Regenerate
Firmware Revision	Optional	SCCRQ	Regenerate
Framing Capabilities	Mandatory	SCCRQ	Regenerate
Framing Type	Mandatory	ICCN	Relay
Host Name	Mandatory	SCCRQ	Regenerate
Message Type	Mandatory	All	Regenerate
Physical Channel Id	Optional	ICRQ	Regenerate
Private Group Id	Optional	ICCN	Relay
Protocol Version	Mandatory	SCCRQ	Regenerate
Proxy Authentication	Optional	ICCN	Relay if acceptable; otherwise regenerate

Copyright © 2013, Juniper Networks, Inc.

Table 44: Default Action for Handling L2TP AVPs at the Switching Boundary *(continued)*

AVP Name	AVP Type	L2TP Message Type	Default Action
Proxy LCP	Optional	ICCN	Relay if acceptable; otherwise regenerate
Receive Window Size	Optional	SCCRQ	Regenerate
Rx Connect Speed	Optional	ICCN	Relay
Sequencing Required	Optional	ICCN	Regenerate
Sub-Address	Optional	ICRQ	Relay
Tie Breaker	Optional	SCCRQ	Regenerate
Tunnel Recovery	Optional	SCCRQ	Regenerate
Tx Connect Speed	Mandatory	ICCN	Relay
Vendor Name	Optional	SCCRQ	Regenerate

Related Documentation

- L2TP Tunnel Switching Overview on page 392
- Configuring L2TP Tunnel Switching on page 432

LAC Interoperation with Third-Party LNS Devices

In some network environments, the LAC may need to interoperate with an LNS configured on a device from another vendor that does not run Junos OS. Interoperation with Cisco Systems devices requires the LAC to communicate a NAS port type, but the LAC does not provide this information by default. You can enable interoperation with Cisco devices by configuring the LAC to send information to the LNS about the NAS port and port type.

You can use either of the following ways to specify the LAC's NAS port method, which controls whether the LAC sends the NAS port information to the LNS:

- Include the **nas-port-method** statement in the tunnel profile configured on the LAC at the **[edit access tunnel-profile]** hierarchy level. Specify **cisco-avp** as the method.
- Include the Tunnel-Nas-Port-Method VSA [26–30] in your RADIUS server configuration with the value set to 1 to indicate Cisco CLID.

The value configured in RADIUS has precedence over the value configured in the CLI when both are configured.

Both methods cause the LAC to include the Cisco Nas Port Info AVP (100) when it sends an incoming call request (ICRQ) to the LNS. The AVP includes information that identifies the NAS port and indicates whether the port type is ATM or Ethernet.

 Copyright © 2013, Juniper Networks, Inc.

When the LNS is a Cisco device, it uses the AVP to facilitate interoperation. When the LNS is an MX Series router, the LNS simply ignores this AVP, unless the LNS is configured for L2TP tunnel switching. In that case, the LNS preserves the value of the AVP and passes it along when it switches tunnels for the LAC.

Related Documentation

- L2TP for Subscriber Access Overview on page 379
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- Configuring an L2TP LAC on page 402

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 16

Configuring L2TP for Subscriber Access

- Configuring an L2TP LAC on page 402
- Configuring a Tunnel Profile for Subscriber Access on page 403
- Configuring the L2TP LAC Tunnel Selection Parameters on page 406
- Configuring LAC Tunnel Selection Failover Within a Preference Level on page 406
- Configuring Weighted Load Balancing for LAC Tunnel Sessions on page 407
- Preventing the LAC from Sending Calling Number AVP 22 to the LNS on page 407
- Configuring the Method to Set the LAC Connection Speeds to the LNS on page 408
- Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal on page 409
- Preventing the LAC From Negotiating L2TP Failover Protocol on page 409
- Setting the Format for the Tunnel Name on page 410
- Configuring the Number of L2TP Control Message Retransmissions on page 411
- Setting the L2TP Tunnel Idle Timeout on page 412
- Setting the L2TP Receive Window Size on page 412
- Setting the L2TP Destruct Timeout on page 413
- Configuring the LAC to Ignore Address and Port Changes Requested by the LNS on page 413
- Configuring the L2TP Destination Lockout Timeout on page 414
- Removing an L2TP Destination from the Destination Lockout List on page 415
- Enabling Tunnel and Global Counters for SNMP Statistics Collection on page 416
- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 418
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 419
- Configuring IP Inline Reassembly for L2TP on page 421
- Configuring an L2TP Access Profile on the LNS on page 422
- Configuring a AAA Local Access Profile on the LNS on page 424
- Configuring an Address-Assignment Pool for L2TP LNS with Inline Services on page 424
- Configuring the L2TP LNS Peer Interface on page 426

Copyright © 2013, Juniper Networks, Inc.

- Enabling Inline Service Interfaces on page 427
- Configuring an Inline Service Interface for L2TP LNS on page 427
- Configuring Options for the LNS Inline Services Logical Interface on page 428
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429
- Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 430
- Configuring a Dynamic Profile for Dynamic LNS Sessions on page 431
- Configuring L2TP Tunnel Switching on page 432
- Tracing L2TP Operations for Subscriber Access on page 434
- Configuring the L2TP Trace Log Filename on page 435
- Configuring the Number and Size of L2TP Log Files on page 435
- Configuring Access to the L2TP Log File on page 436
- Configuring a Regular Expression for L2TP Messages to Be Logged on page 436
- Configuring the L2TP Tracing Flags on page 437
- Configuring the Severity Level to Filter Which L2TP Messages Are Logged on page 437
- Verifying and Managing L2TP for Subscriber Access on page 438
- Testing L2TP Tunnel Configurations from the LAC on page 439
- Example: Configuring an L2TP LNS on page 440

Configuring an L2TP LAC

To configure an L2TP LAC:

1. Configure a tunnel profile to apply to subscribers.

   See "Configuring a Tunnel Profile for Subscriber Access" on page 403.

2. (Optional) Configure the method used for selecting among multiple tunnels.

   - See "Configuring the L2TP LAC Tunnel Selection Parameters" on page 406.
   - See "Configuring Weighted Load Balancing for LAC Tunnel Sessions" on page 407.
   - See "Configuring LAC Tunnel Selection Failover Within a Preference Level" on page 406.

3. (Optional) Configure the LAC to not send Calling Number AVP 22 to the LNS.

   See "Preventing the LAC from Sending Calling Number AVP 22 to the LNS" on page 407.

4. (Optional) Specify the method for setting the transmit and receive connect speeds.

   See "Configuring the Method to Set the LAC Connection Speeds to the LNS" on page 408.

5. (Optional) Disable negotiation of the L2TP failover protocol to force use of only the silent failover resynchronization mechanism.

   See "Preventing the LAC From Negotiating L2TP Failover Protocol" on page 409.

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

 Copyright © 2013, Juniper Networks, Inc.

"Setting the Format for the Tunnel Name" on page 410.

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

   See "Configuring the Number of L2TP Control Message Retransmissions" on page 411.

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

   See "Setting the L2TP Tunnel Idle Timeout" on page 412.

9. (Optional) Specify the L2TP receive window size for the L2TP tunnel. The receive window size specifies the number of packets a peer can send before waiting for an acknowledgment from the router.

   See "Setting the L2TP Receive Window Size" on page 412

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

    See "Setting the L2TP Destruct Timeout" on page 413.

11. (Optional) Specify how the LAC handles IP address or UDP port change requests.

    See "Configuring the LAC to Ignore Address and Port Changes Requested by the LNS" on page 413

12. (Optional) Enable SNMP statistics counters.

    See "Enabling Tunnel and Global Counters for SNMP Statistics Collection" on page 416

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

    See "Tracing L2TP Operations for Subscriber Access" on page 434

Configuring a Tunnel Profile for Subscriber Access

The tunnel profile specifies a set of attributes to characterize the tunnel. The profile can be applied by a domain map or automatically when the tunnel is created.

NOTE: RADIUS attributes and VSAs can override the values you configured by a tunnel profile in a domain map. In the absence of a domain map, RADIUS can supply all the characteristics of a tunnel. The steps in the following procedure list the corresponding standard RADIUS attribute or VSA that you can configure on your RADIUS server to modify or configure the tunnel profile.

RADIUS-supplied attributes are associated with a tunnel by a tag carried in the attribute, which matches the tunnel identifier. A tag of 0 indicates the tag is not used. If L2TP receives a RADIUS attribute with a tag of 0, the attribute cannot be merged with the tunnel profile configuration corresponding to the subscriber domain because a tunnel profile cannot provide a tunnel tag (tunnel identifier) of 0. Only tags in the range of 1 through 31 are supported.

To configure a tunnel definition for a tunnel profile:

1. Specify the tunnel profile for which you are defining a tunnel. (Tunnel-Group [26-64])

   [edit access]
   user@host# set tunnel-profile *profile-name*

2. Specify an identifier (name) for the L2TP control connection for the tunnel.

   [edit access tunnel-profile *profile-name*]
   user@host# set tunnel *tunnel-id*

3. Configure the IP address of the local L2TP tunnel endpoint, the LAC. (Tunnel-Client-Endpoint [66])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set source-gateway address *client-ip-address*

4. Configure the IP address of the remote L2TP tunnel endpoint, the LNS. (Tunnel-Server-Endpoint [67])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set remote-gateway address *server-ip-address*

5. (Optional) Configure the preference level for the tunnel. (Tunnel-Preference [83])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set preference *number*

6. (Optional) Configure the hostname of the local client (LAC). (Tunnel-Client-Auth-Id [90])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set source-gateway gateway-name *client-name*

7. (Optional) Configure the hostname of the remote server (LNS). (Tunnel-Server-Auth-Id [91])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set remote-gateway gateway-name *server-name*

8. (Optional) Specify the medium (network) type for the tunnel. (Tunnel-Medium-Type [65])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set medium *type*

9. (Optional) Specify the protocol type for the tunnel. (Tunnel-Type [64])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set type *tunnel-type*

10. (Optional) Configure the assignment ID for the tunnel. (Tunnel-Assignment-Id [82])

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set identification *name*

11. (Optional) Configure the maximum number of sessions allowed in the tunnel. (Tunnel-Max-Sessions [26-33])

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set max-sessions *number*

Copyright © 2013, Juniper Networks, Inc.

12. (Optional) Configure the password for remote server authentication. (Standard RADIUS attribute Tunnel-Password [69] or VSA Tunnel-Password [26-9])

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set secret ***password***

13. (Optional) Configure the logical system to use for the tunnel.

    If you configure a logical system, you must also configure a routing instance.

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set logical-system ***logical-system-name***

14. (Optional) Configure the routing instance to use for the tunnel. (Tunnel-Virtual-Router [26-8])

    If you configure a routing instance, configuring a logical system is optional.

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set routing-instance ***routing-instance-name***

15. (Optional) Enable the LAC to interoperate with Cisco LNS devices.

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set nas-port-method cisco-avp

The following example shows a complete configuration for a tunnel profile:

```
tunnel-profile marketing {
  tunnel 1 {
    preference 5;
    remote-gateway {
      address 172.16.98.4;
      gateway-name work;
    }
    source-gateway {
      address 192.168.4.10;
      gateway-name local;
    }
    secret mk5Sn$3k%V;
    logical-system bos-metro-5;
    routing-instance rox-12-32;
    medium ipv4;
    type l2tp;
    identification tunnel_to_work;
    max-sessions 32;
    nas-port-method cisco avp;
  }
}
```

Related Documentation

- Configuring an L2TP LAC on page 402
- Domain Mapping Overview on page 180
- LAC Interoperation with Third-Party LNS Devices on page 398

Configuring the L2TP LAC Tunnel Selection Parameters

When the LAC determines that a PPP session should be tunneled, it selects a tunnel from the set of tunnels associated with either the PPP user or the PPP user's domain. You can configure how a tunnel is selected and whether certain information is sent by the LAC to the LNS.

To configure tunnel selection parameters:

1. (Optional) Configure how a tunnel is selected when a connection attempt fails.

   See "Configuring LAC Tunnel Selection Failover Within a Preference Level" on page 406.

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

   See "Configuring Weighted Load Balancing for LAC Tunnel Sessions" on page 407.

Related Documentation

- LAC Tunnel Selection Overview on page 384

Configuring LAC Tunnel Selection Failover Within a Preference Level

You can configure how LAC tunnel selection continues in the event of a failure to connect. By default, when the router is unable to connect to a destination at a given preference level, it attempts to connect at the next lower level. You can specify that the router instead attempt to connect to another destination at the same level as the failed attempt.

If all destinations at a preference level are marked as unreachable, the router does not attempt to connect to a destination at that level. It drops to the next lower preference level to select a destination.

If all destinations at all preference levels are marked as unreachable, the router chooses the destination that failed first and tries to make a connection. If the connection fails, the router rejects the PPP user session without attempting to contact the remote router.

For example, suppose there are four tunnels for a domain: A, B, C, and D. All tunnels are considered reachable, and the preference levels are assigned as follows:

- A and B at preference 0
- C and D at preference 1

When the router attempts to connect to the domain, suppose it randomly selects tunnel B from preference 0. If it fails to connect to tunnel B, the router excludes tunnel B for five minutes and attempts to connect to tunnel A. If this attempt also fails, the router drops to preference 1. Then suppose the router selects tunnel C. If it also fails to connect to tunnel C, the router excludes tunnel C for five minutes and attempts to connect to tunnel D.

You configure the preference level used for this tunnel selection method in the tunnel profile or the RADIUS Tunnel-Preference [83] attribute.

 Copyright © 2013, Juniper Networks, Inc.

To enable tunnel selection failover within a preference level:

- Specify failover within preference.

```
[edit services l2tp]
user@host# set fail-over-within-preference
```

Related Documentation

- LAC Tunnel Selection Overview on page 384
- Configuring the L2TP LAC Tunnel Selection Parameters on page 406
- Configuring a Tunnel Profile for Subscriber Access on page 403
- Configuring How RADIUS Attributes Are Used for Subscriber Access on page 58

Configuring Weighted Load Balancing for LAC Tunnel Sessions

You can configure how L2TP LAC sessions are distributed across tunnels. You can specify that the router uses the maximum sessions per tunnel to choose among multiple tunnels that share the same preference level.

The weight of a tunnel is proportional to its maximum session limit and the maximum session limits of the other tunnels at the same preference level. The tunnel with the largest maximum session value has the highest weight. The tunnel with the next larger maximum session value has the next higher weight, and so on. The tunnel with the smallest maximum session value has the lowest weight.

When you configure weighted load balancing, the tunnel with the highest weight in the preference level is selected until the maximum number of sessions for the tunnel is reached. Then the router selects the tunnel with the next higher weight to establish connections until that tunnel's maximum session limit is reached, and so on.

To configure weighted load balancing:

- Specify load balancing.

```
[edit services l2tp]
user@host# set weighted-load-balancing
```

Related Documentation

- LAC Tunnel Selection Overview on page 384
- Configuring the L2TP LAC Tunnel Selection Parameters on page 406

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

Calling Number AVP 22 typically identifies the interface that is connected to the customer in the access network. When RADIUS includes the Calling-Station-Id in the Access-Accept message, that value is used for the Calling Number AVP. Otherwise, the underlying interface (for example, the S-VLAN IFL) on which the PPPoE session is established is used for the Calling Number AVP value.

Copyright © 2013, Juniper Networks, Inc.

By default, the LAC includes this AVP in the incoming-call request (ICRQ) packets that it sends to the LNS. However, you may wish to hide your network access interface information. To do so, you can configure the tunnel so that the LAC does not send the Calling Number AVP to the LNS.

To disable sending the Calling Number AVP:

- Configure disabling.

```
[edit services l2tp]
user@host# set disable-calling-number-avp
```

Related Documentation

- LAC Tunnel Selection Overview on page 384

Configuring the Method to Set the LAC Connection Speeds to the LNS

During the establishment of an L2TP tunnel session, the LAC sends the L2TP transmit connect speed in bits per second (BPS) using AVP 24 to the LNS in Incoming-Call-Connected (ICCN) messages. AVP 24 conveys the transmit connect speed of the subscriber's access interface; that is, it represents the speed of the connection from the LAC to the LNS, from the perspective of the LAC. The L2TP receive connect speed, which is represented by AVP 38, is included in the message when the receive connect speed is different from the transmit connect speed. AVP 38 conveys the receive connect speed of the connection from the LNS to the LAC, again from the perspective of the LAC. When AVP 38 is not sent, it means that the connection speed is the same in both directions; the LNS uses the value in AVP 24 for both transmit and receive connect speeds.

You can configure what the LAC uses as a resource for setting these speeds. To use the recommended (advisory) downstream shaping rate for AVP 24 and the recommended upstream shaping rate for AVP 38, include the **tx-connect-speed-method static** statement at the **[edit services l2tp]** hierarchy level. You configure the advisory rates under the PPPoE logical interface underlying the subscriber interface with the **advisory-options** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level. When the advisory speed is not configured on the underlying interface, the **tx-connect-speed-method advisory** statement automatically sets the speed to 1 Gbps and sends this value in both AVP 24 and AVP 38.

To derive the speeds from the PPPoE IA tags, use the **tx-connect-speed-method pppoe-ia-tags** statement. In this case, AVP 24 is the value of Actual-Data-Rate-Downstream (VSA 26-129). AVP 38 is the value of Actual-Data-Rate-Upstream (VSA 26-130), and is sent only when the VSA values differ.

To derive the speeds from the ANCP value configured on the PPPoE interface underlying the subscriber interface, use the **tx-connect-speed-method ancp** statement.

To set the method for calculating the transmit connect speed:

- Configure the ANCP method to use the values derived from the configured ANCP value for the underlying interface.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit services l2tp]
user@host# set tx-connect-speed-method ancp
```

- Configure the PPPoE IA tags method to use the values provided in the PPPoE IA tags.

```
[edit services l2tp]
user@host# set tx-connect-speed-method pppoe-ia-tags
```

- Configure the static (advisory downstream shaping rate) method to use the underlying interface's recommended shaping rates.

```
[edit services l2tp]
user@host# set tx-connect-speed-method static
```

Related Documentation

- Configuring an L2TP LAC on page 402

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

The L2TP Rx Connect Speed (in bits per second) AVP, which is represented by AVP 38, is included in the ICCN message when the receive connect speed is different from the transmit connect speed. By default, when the connection speed is the same in both directions, AVP 38 is not sent; the LNS uses the value in AVP 24 for both transmit and receive connect speeds.

AVP 38 is generated when the receive connect speed of the access interface is set equal to the calculated transmit connect speed by issuing the **rx-connect-speed-when-equal** statement at the **[edit services l2tp]** hierarchy level. In this scenario, the LAC transmits the same value for transmit and receive connect speeds that are sent to the LNS through the AVP 24 and AVP 38 in the ICCN message.

To configure the sending of AVP 38 when the connection speeds are the same in both the downstream and upstream directions:

- Configure the transmission of the receive connect speed, AVP 38, when the receive connect speed is set equal to the calculated transmit connect speed.

```
[edit services l2tp]
user@host# set rx-connect-speed-when-equal
```

Related Documentation

- Transmission of Tx Connect-Speed and Rx Connect-Speed AVPs from LAC to LNS on page 391
- Configuring an L2TP LAC on page 402
- **rx-connect-speed-when-equal on page 1967**

Preventing the LAC From Negotiating L2TP Failover Protocol

The L2TP LAC implementation on MX Series routers supports L2TP failover and peer resynchronization with a failed remote endpoint. The LAC supports both the L2TP failover protocol method and the L2TP silent failover method. By default, L2TP on the LAC

attempts to negotiate the L2TP failover protocol with the LNS. When negotiation determines that the LNS supports this method, then the LAC uses L2TP failover protocol if the LNS fails. When the LNS does not support L2TP failover protocol, then the LAC uses silent failover in the event of an LNS failure. The ability to fall back on silent failover prevents the failover from forcing a disconnection of the tunnel to the peer and all the associated sessions.

You can disable the default behavior to force the LAC to operate only in silent failover mode. This configuration can be useful when routers that act as the LNS either are configured for silent failover or incorrectly negotiate use of the failover protocol even though they do not support it. However, when you issue this statement and the LNS supports only failover protocol, then the LAC cannot negotiate failover protocol, and recovery (failover protocol recovery initiated by the LNS) always fails.

To disable negotiation of the L2TP failover protocol:

- Configure disabling.

```
[edit services l2tp]
user@host# set disable-failover-protocol
```

Related Documentation

- Configuring an L2TP LAC on page 402

Setting the Format for the Tunnel Name

By default, the name of a tunnel corresponds to the Tunnel-Assignment-Id [82] returned by the AAA server. You can optionally configure the LAC to use more elements in the construction of a tunnel name by including the **assignment-id-format client-server-id** statement at the **[edit services l2tp tunnel]** hierarchy level. This format uses three attributes: Tunnel-Client-Auth-Id [90], Tunnel-Server-Endpoint [67], and Tunnel-Assignment-Id [82]. These attributes correspond, respectively, to the values configured in the tunnel profile for the LAC (source gateway) name, the tunnel endpoint (remote gateway) address on the LNS, and the tunnel ID.

A consequence of the **client-server-id** format is that the LAC automatically creates a new tunnel when the AAA server returns a different Tunnel-Client-Auth-Id than previously returned.

NOTE: Before you downgrade to a Junos OS Release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the **no assignment-id-format assignment-id** statement at the **[edit services l2tp tunnel]** hierarchy level.

To change how the tunnel name is formatted:

- Configure the format.

```
[edit services l2tp tunnel]
user@host# set assignment-id-format client-server-id
```


Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring an L2TP LAC on page 402

Configuring the Number of L2TP Control Message Retransmissions

L2TP peers maintain a queue of control messages that must be sent to the peer device. After a message is sent, the local peer waits for a response from the remote peer. If a response is not received, the local peer retransmits the message. You can configure how many times an unacknowledged message is retransmitted by the LAC or the LNS. For tunnels that have been established, include the **retransmission-count-established** statement at the **[edit services l2tp tunnel]** hierarchy level. For tunnels that are not yet established, include the **retransmission-count-not-established** statement.

The local peer waits one second for the first response to a control message. The retransmit timer then doubles the interval between each successive retransmission, up to a maximum interval of 16 seconds. This behavior allows the remote peer more time to respond. If the maximum retransmission count is reached and no response has been received, the tunnel and all its sessions are cleared.

BEST PRACTICE: Before you downgrade to a Junos OS release that does not support these statements, we recommend that you explicitly unconfigure the feature by including the **no retransmission-count-established** statement and the **no retransmission-count-non-established** statement at the **[edit services l2tp tunnel]** hierarchy level.

BEST PRACTICE: During a unified in-service software upgrade (unified ISSU) on an MX Series router configured as the LAC, the LAC does not respond to control messages from the LNS. This can result in dropping LAC L2TP sessions. You can avoid this situation by ensuring that the maximum retransmission count on the LNS is set to 16 or higher.

To set the maximum retransmission count for established tunnels:

- Configure the count.

```
[edit services l2tp tunnel]
user@host# set retransmission-count-established count
```

To set the maximum retransmission count for non-established tunnels:

- Configure the count.

```
[edit services l2tp tunnel]
user@host# set retransmission-count-not-established count
```

Related Documentation

- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Setting the L2TP Tunnel Idle Timeout

You can configure the LAC or the LNS to specify how long a tunnel without any sessions remains active. The idle timer starts when the last session on the tunnel is terminated. When the timer expires the tunnel is disconnected. This idle timeout frees up resources otherwise consumed by inactive tunnels.

If you set the idle timeout value to zero, the tunnel is forced to remain active indefinitely after the last session is terminated until one of the following occurs:

- You issue the **clear services l2tp tunnel** command.
- The remote peer disconnects the tunnel.

BEST PRACTICE: **Before you downgrade to a Junos OS Release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the no idle-timeout statement at the [edit services l2tp tunnel] hierarchy level.**

To set the tunnel idle timeout:

- Configure the timeout period.

```
[edit services l2tp tunnel]
user@host# set idle-timeout seconds
```

Related Documentation

- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Setting the L2TP Receive Window Size

You can configure the L2TP receive window size for an L2TP tunnel. The receive window size specifies the number of packets a peer can send before waiting for an acknowledgment from the router.

By default, the receive window size is set to four packets. If the receive window size is set to its default value, the router does not send the Receive Window Size AVP, AVP 10, in its first packet sent during tunnel negotiation to its peer.

To configure the receive window size:

```
[edit services l2tp tunnel]
user@host# set rx-window-size packets
```

Related Documentation

- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Copyright © 2013, Juniper Networks, Inc.

Setting the L2TP Destruct Timeout

You can configure the LAC or the LNS to specify how long the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed. This destruct timeout aids debugging and other analysis by saving underlying memory structures after the destination, tunnel, or session is terminated. Any specific dynamic destination, tunnel, or session may not be maintained for this entire time period if the resources must be reclaimed early to allow new tunnels to be established.

BEST PRACTICE: **Before you downgrade to a Junos OS Release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the no destruct-timeout statement at the [edit services l2tp] hierarchy level.**

To set the L2TP destruct timeout:

- Configure the timeout period.

  [edit services l2tp]
  user@host# set destruct-timeout *seconds*

Related Documentation

- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Configuring the LAC to Ignore Address and Port Changes Requested by the LNS

By default, when the LAC receives a request from the LNS in an SCCRP message to change the destination IP address or the UDP port, the LAC accepts the request and makes the change. If this is not the desired behavior, you can use the **tx-address-change** statement to configure how the LAC handles these change requests.

To configure how the LAC handles change requests for the IP address, the UDP port, or both:

- (Optional) Specify the LAC to ignore all change requests.

  [edit services l2tp tunnel]
  user@host# set tx-address-change ignore

- (Optional) Specify the LAC to ignore only change requests for the IP address.

  [edit services l2tp tunnel]
  user@host# set tx-address-change ignore-ip-address

- (Optional) Specify the LAC to ignore only change requests for the UDP port.

  [edit services l2tp tunnel]
  user@host# set tx-address-change ignore-udp-port

- (Optional) Specify the LAC to accept all change requests.

[edit services l2tp tunnel]
user@host# set tx-address-change accept

Use the **show services l2tp summary** command to display the current behavior of the LAC:

```
show services l2tp summary
  Failover within a preference level is Disabled
  Weighted load balancing is Disabled
  Tunnel authentication challenge is Enabled
  Calling number avp is Enabled
  Failover Protocol is Disabled
  Tx Connect speed method is static
  Rx speed avp when equal is Disabled
  Tunnel assignment id format is assignment-id
  Tunnel Tx Address Change is Ignore IP Address Change
  Max Retransmissions for Established Tunnel is 7
  Max Retransmissions for Not Established Tunnel is 5
  Tunnel Idle Timeout is 60 seconds
  Destruct Timeout is 300 seconds
  Destination Lockout Timeout is 300 seconds
Destinations: 1, Tunnels: 0, Sessions: 0
```

Related Documentation

- Configuring an L2TP LAC on page 402

Configuring the L2TP Destination Lockout Timeout

When multiple sets of tunneling parameters are available, L2TP uses a selection process to choose the best tunnel for subscriber traffic. As part of this selection process, L2TP locks out destinations it cannot connect to when a subscriber tries to reach a domain. L2TP places the destination on the destination lockout list and excludes the destination from consideration for a configurable period called the *destination lockout timeout*.

By default, the destination lockout timeout is 300 seconds (5 minutes). You can configure a value from 60 through 3600 seconds (1 minute through 1 hour). When the lockout timeout expires, L2TP assumes that the destination is now available and includes the destination when performing the tunnel selection process. The destination lockout period is a global value and is not individually configurable for particular destinations, tunnels, or tunnel groups.

BEST PRACTICE: **Configure the lockout timeout to be equal to or shorter than the destruct timeout. Otherwise, the destruct timeout expires before the lockout timeout. In this event, the locked-out destination is destroyed and can be subsequently returned to service before the lockout timeout expires, thus negating the effectiveness of the lockout timeout.**

To configure the destination lockout timeout:

- Specify the period in seconds.

  [edit services l2tp destination]
  user@host# set lockout-timeout *seconds*

Copyright © 2013, Juniper Networks, Inc.

The **show services l2tp destination lockout** command displays the destination lockout list and for each destination indicates how much time remains before its timeout expires. The **show services l2tp destination detail** command indicates for each destination whether it is locked and waiting for the timeout to expire or not locked.

Related Documentation

- LAC Tunnel Selection Overview on page 384
- Setting the L2TP Destruct Timeout on page 413
- Removing an L2TP Destination from the Destination Lockout List on page 415
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Removing an L2TP Destination from the Destination Lockout List

When a PPP subscriber tries to log in to a domain, L2TP selects a tunnel associated with a destination in that domain and attempts to access the destination. If the connection attempt fails, L2TP places the destination on the destination lockout list. Destinations on this list are excluded from being considered for subsequent connections for a configurable period called the *destination lockout timeout*.

You can issue the **request services l2tp destination unlock** command for a particular destination to remove it from the destination lockout list. The result is that this destination is immediately available for consideration when a subscriber logs in to the associated domain.

To remove a destination from the destination lockout list:

- Specify the name of the destination to be unlocked.

  user@host> **request services l2tp destination unlock** ***destination-name***

Related Documentation

- LAC Tunnel Selection Overview on page 384
- Configuring the L2TP Destination Lockout Timeout on page 414

Enabling Tunnel and Global Counters for SNMP Statistics Collection

By default, SNMP polling is disabled for L2TP statistics. As a consequence, the L2TP tunnel and global counters listed in Table 45 on page 416 have a default value of zero.

Table 45: SNMP Counters for L2TP Statistics

Counter Name	Type
jnxL2tpTunnelStatsDataTxPkts	Tunnel
jnxL2tpTunnelStatsDataRxPkts	Tunnel
jnxL2tpTunnelStatsDataTxBytes	tunnel
jnxL2tpTunnelStatsDataRxBytes	Tunnel
jnxL2tpStatsPayloadRxOctets	Global
jnxL2tpStatsPayloadRxPkts	Global
jnxL2tpStatsPayloadTxOctets	Global
jnxL2tpStatsPayloadTxPkts	Global

You can enable collection of these statistics by including the **enable-snmp-tunnel-statistics** statement at the **[edit services l2tp]** hierarchy level. When enabled, the L2TP process polls for these statistics every 30 seconds for 1000 sessions. The potential age of the statistics increases with the number of subscriber sessions; the data is refreshed more quickly as the number of sessions decreases. For example, with 60,000 sessions, none of these statistics can be more than 30 minutes old.

BEST PRACTICE: **The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.**

To enable L2TP statistics collection for SNMP:

- Enable statistics collection.

```
[edit services l2tp]
user@host1# set enable-snmp-tunnel-statistics
```

Related Documentation

- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Copyright © 2013, Juniper Networks, Inc.

Configuring an L2TP LNS with Inline Service Interfaces

The L2TP LNS feature license must be installed before you begin the configuration. Otherwise, a warning message is displayed when the configuration is committed.

To configure an L2TP LNS with inline service interfaces:

1. (Optional) Configure a user group profile that defines the PPP configuration for tunnel subscribers.

   See "Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile" on page 418.

2. (Optional) Configure PPP attributes for subscribers on inline service interfaces.

   See "Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface" on page 419.

3. Configure an L2TP access profile that defines the L2TP parameters for each LNS client (LAC).

   See "Configuring an L2TP Access Profile on the LNS" on page 422.

4. (Optional) Configure a AAA access profile to override the access profile configured under the routing instance.

   See "Configuring a AAA Local Access Profile on the LNS" on page 424.

5. Configure a pool of addresses to be dynamically assigned to tunneled PPP subscribers.

   See "Configuring an Address-Assignment Pool for L2TP LNS with Inline Services" on page 424.

6. Configure the peer interface to terminate the tunnel and the PPP server-side IPCP address.

   See "Configuring the L2TP LNS Peer Interface" on page 426.

7. Enable inline service interfaces on an MPC.

   See "Enabling Inline Service Interfaces" on page 427.

8. Configure a service interface.

   See "Configuring an Inline Service Interface for L2TP LNS" on page 427.

9. Configure options for each inline service logical interface.

   See "Configuring Options for the LNS Inline Services Logical Interface" on page 428.

10. Configure the L2TP tunnel group.

    See "Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces" on page 429.

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

    See "Configuring a Dynamic Profile for Dynamic LNS Sessions" on page 431

Copyright © 2013, Juniper Networks, Inc.

12. (Optional) Configure a service interface pool for dynamic LNS sessions.

    See "Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions" on page 430.

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

    See "Configuring the Number of L2TP Control Message Retransmissions" on page 411.

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

    See "Setting the L2TP Tunnel Idle Timeout" on page 412.

15. (Optional) Specify the L2TP receive window size for the L2TP tunnel. The receive window size specifies the number of packets a peer can send before waiting for an acknowledgment from the router.

    See "Setting the L2TP Receive Window Size" on page 412

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

    See "Setting the L2TP Destruct Timeout" on page 413.

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

    See "Configuring the L2TP Destination Lockout Timeout" on page 414.

18. (Optional) Configure L2TP tunnel switching.

    See "Configuring L2TP Tunnel Switching" on page 432

19. (Optional) Enable SNMP statistics counters.

    See "Enabling Tunnel and Global Counters for SNMP Statistics Collection" on page 416

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

    See "Tracing L2TP Operations for Subscriber Access" on page 434

You also need to configure CoS for LNS sessions. For more information, see "Configuring Dynamic CoS for an L2TP LNS Inline Service" on page 1026.

Related Documentation

- L2TP for Subscriber Access Overview on page 379
- *Junos OS Feature Licenses*
- *Software Feature Licenses*

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

You can configure a user group profile that enables the LNS to apply PPP attributes to the PPP subscribers tunneled from the LAC. The user group profile is associated with clients (LACs) in the L2TP access profile. Consequently all subscribers handled by a given client share the same PPP attributes.

 Copyright © 2013, Juniper Networks, Inc.

To configure a user group profile:

1. Create the profile.

   [edit access]
   user@host# **edit group-profile** *profile-name*

2. Configure the interval between PPP keepalive messages for the L2TP tunnel terminating on the LNS.

   [edit access group-profile *profile-name*]
   user@host# **set ppp** keepalive *seconds*

   NOTE: Changes to the keepalive interval in a user group profile affect only new L2TP sessions that come up after the change. Existing sessions are not affected.

3. Configure PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

   [edit access group-profile *profile-name*]
   user@host# **set ppp** ppp-options chap
   user@host# **set ppp** ppp-options pap

4. Configure how long the PPP subscriber session can be idle before it is considered to have timed out.

   [edit access group-profile *profile-name*]
   user@host# **set ppp** idle-timeout 200

   NOTE: You can also configure PPP attributes on a per-interface basis. See "Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface" on page 419 for more information. When you configure the PPP attributes for L2TP LNS subscribers both on the si interface and in user group profiles, the inline service interface configuration takes precedence over the user group profile configuration.

Related Documentation

- Configuring an L2TP Access Profile on the LNS on page 422
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface

You can configure PPP attributes that are applied by the LNS on the inline service (si) interface to the PPP subscribers tunneled from the LAC. Because you are configuring the attributes per interface rather than with a user group profile, the attributes for subscribers can be varied with a finer granularity. This configuration matches that used for terminated PPPoE subscribers.

To configure the PPP attributes for dynamically created si interfaces:

1. Specify the predefined dynamic interface and logical interface variables in the dynamic profile.

   [edit dynamic-profiles *profile-name*]
   user@host# **edit interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"**

2. Configure the interval between PPP keepalive messages for the L2TP tunnel terminating on the LNS.

   [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"]
   user@host# **set keepalives interval *seconds***

3. Configure PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

   [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"]
   user@host# **set ppp-options chap**
   user@host# **set ppp-options pap**

To configure the PPP attributes for statically created si interfaces:

1. Specify the logical inline service interface.

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

2. Configure the interval between PPP keepalive messages for the L2TP tunnel terminating on the LNS.

   [edit interfaces si-*slot/pic/port* unit *logical-unit-number*]
   user@host# **set keepalives interval *seconds***

3. Configure the number of keepalive packets a destination must fail to receive before the network takes down a link.

   [edit interfaces si-*slot/pic/port* unit *logical-unit-number*]
   user@host# **set keepalives down-count *number***

   NOTE: **The keepalives up-count option is typically not used for subscriber management.**

4. Configure PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

   [edit interfaces si-*slot/pic/port* unit *logical-unit-number*]
   user@host# **set ppp-options chap**
   user@host# **set ppp-options pap**

   BEST PRACTICE: **Although all other statements subordinate to ppp-options—including those subordinate to chap and pap—are supported,**

Copyright © 2013, Juniper Networks, Inc.

they are typically not used for subscriber management. We recommend that you leave these other statements at their default values.

NOTE: You can also configure PPP attributes with a user group profile that applies the attributes to all subscribers with that profile on a LAC client. See "Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile" on page 418 for more information. When you configure the PPP attributes for L2TP LNS subscribers both on the si interface and in user group profiles, the inline service interface configuration takes precedence over the user group profile configuration.

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Configuring IP Inline Reassembly for L2TP

This procedure shows how to configure a service interface on a LAC or LNS to reassembly fragmented IP packets. This example creates a service set that configures the IP reassembly parameters for L2TP fragments. The service set is then associated with the L2TP service.

Before you configure inline IP reassembly, be sure you have:

- Configured L2TP.
- Configured a valid service interface on the LAC or LNS.

To configure inline IP reassembly:

1. Configure the chassis-level bandwidth used by the inline services interface on the FPC and PIC slot for inline IP fragment reassembly.

   ```
   [edit chassis]
   user@host# set fpc 2 pic 1 inline-services bandwidth 10g
   ```

2. Configure the interface-level logical unit used by the inline services (**si-**) interface on the FPC and PIC slot for inline IP fragment reassembly.

   ```
   [edit interfaces]
   user@host# set si-2/1/0 unit 0 family inet
   user@host# set si-2/1/0 unit 0 service-domain inside
   ```

   NOTE: This configuration is not unique to L2TP. However, you must configure the family (**inet**) and service domain (**inside**) as shown.

3. Configure the service set (**set1**) for IP reassembly in the input match direction. (The **local** option loops the reassembled packets back to the local interface.)

   ```
   [edit services]
   ```

```
user@host# set service-set set1

[edit services service set ip-reassembly-set]
user@host# set ip-reassembly-rules ipr_rule1
user@host# set next-hop-service inside-service-interface si-9/1/0.0
user@host# set next-hop-service outside-service-interface-type local
```

NOTE: You must configure both inside (si- interface) and outside type (local) service interfaces statements. The reassembly rule is not formulated outside of the service set; this statement simply initiates the reassembly process.

4. Configure the IP reassembly rule parameter

```
[edit services ip-reassembly]
user@host# set rule ipr_rule1 match-direction input;
```

5. Configure the service set (**set1**) for IP reassembly to bind to the L2TP service.

NOTE:
- The service set must be defined at the **[edit services]** hierarchy level.
- You cannot delete a service set instance if it is associated with an L2TP service.

```
[edit services l2tp]
user@host# set ip-reassembly service-set set1
```

Related Documentation

- IP Packet Fragment Reassembly for L2TP Overview on page 390
- *Protocols and Applications Supported by MX240, MX480, MX960, and MX2020 Enhanced MPCs (MPCEs)*
- ip-reassembly on page 1742

Configuring an L2TP Access Profile on the LNS

Access profiles define how to validate Layer 2 Tunneling Protocol (L2TP) connections and session requests. Within each L2TP access profile, you configure one or more clients (LACs). The client characteristics are used to authenticate LACs with matching passwords, and to establish attributes of the client tunnel and session. You can configure multiple access profiles and multiple clients within each profile.

To configure an L2TP access profile:

1. Create the access profile.

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

2. Configure characteristics for one or more clients (LACs).

 Copyright © 2013, Juniper Networks, Inc.

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

NOTE: Except for the special case of the **default** client, the LAC client name that you configure in the access profile must match the hostname of the LAC. In the case of a Juniper Networks router acting as the LAC, the hostname is configured in the LAC tunnel profile with the gateway gateway-name statement at the **[edit access tunnel-profile *profile-name* tunnel *tunnel-id* source-gateway]** hierarchy level. Alternatively, the client name can be returned from RADIUS in the attribute, Tunnel-Client-Auth-Id [90].

NOTE: Use **default** as the client name when you want to define a default tunnel client. The default client enables the authentication of multiple LACs with the same secret and L2TP attributes. This behavior is useful when, for example, many new LACs are added to the network, because it enables the LACs to be used without additional LNS profile configuration.

Use **default** only on MX Series routers. The equivalent client name on M Series routers is *.

3. (Optional) Specify a local access profile that overrides the global access profile and the tunnel group AAA access profile to configure RADIUS server settings for the client.

```
[edit access profile access-profile-name client client-name]
user@host# set l2tp aaa-access-profile
```

4. Configure the LNS to renegotiate the link control protocol (LCP) with the PPP client. tunneled from the client.

```
[edit access profile access-profile-name client client-name]
user@host# set l2tp lcp-renegotiation
```

5. Configure the maximum number of sessions allowed in a tunnel from the client (LAC).

```
[edit access profile access-profile-name client client-name]
user@host# set l2tp maximum-sessions-per-tunnel number
```

6. Configure the tunnel password used to authenticate the client (LAC).

```
[edit access profile access-profile-name client client-name]
user@host# set l2tp shared-secret shared-secret
```

7. (Optional) Associate a group profile containing PPP attributes to apply for the PPP sessions being tunneled from this LAC client.

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

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429

Configuring a AAA Local Access Profile on the LNS

For some LNS tunnels, you might wish to override the access profile configured at the routing instance that hosts the tunnel with a particular RADIUS server configuration. You can configure a local access profile to do so. You can subsequently use the **aaa-access-profile** statement to apply the local access profile to a tunnel group or LAC client.

A local access profile applied to a client overrides a local access profile applied to a tunnel group, which in turn overrides the access profile for the routing instance.

To configure a AAA local access profile:

1. Create the access profile.

   [edit access]
   user@host# **edit** profile ***local-aaa-profile-name***

2. Configure the order of AAA authentication methods.

   [edit access profile *local-aaa-profile-name*]
   user@host# **set** authentication-order **radius**

3. Configure the RADIUS server attributes, such as the authentication password.

   [edit access profile *local-aaa-profile-name*]
   user@host# **set** radius-server ***server-address*** secret ***password***

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429
- Configuring an L2TP Access Profile on the LNS on page 422

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

You can configure pools of addresses that can be dynamically assigned to the tunneled PPP subscribers. The pools must be local to the routing instance where the subscriber comes up. The configured pools are supplied in the RADIUS Framed-Pool and Framed-IPv6-Pool attributes. Pools are optional when Framed-IP-Address is sent by RADIUS.

To configure an address-assignment pool, you must specify the name of the pool and configure the addresses for the pool.

You can optionally configure multiple named ranges, or subsets, of addresses within an address-assignment pool. During dynamic address assignment, a client can be assigned an address from a specific named range. To create a named range, you specify a name for the range and define the address range.

Copyright © 2013, Juniper Networks, Inc.

NOTE: Be sure to use the address-assignment pools (address-assignment) statement rather than the address pools (address-pool) statement.

To configure an IPv4 address-assignment pool for L2TP LNS:

1. Configure the name of the pool and specify the IPv4 family.

   [edit access]
   user@host# edit address-assignment pool *pool-name* family inet

2. Configure the network address and the prefix length of the addresses in the pool.

   [edit access address-assignment pool *pool-name* family inet]
   user@host# set network *ip-prefix/<prefix-length*

3. Configure the name of the range and the lower and upper boundaries of the addresses in the range.

   [edit access address-assignment pool *pool-name* family inet]
   user@host# set range *range-name* low *lower-limit* high *upper-limit*

For example, to configure an IPv4 address-assignment pool:

```
[edit access]
user@host# edit address-assignment pool lns-v4-pool family inet
[edit access address-assignment pool lns-v4-pool family inet]
user@host# set network 192.168.1.1/16
[edit access address-assignment pool lns-v4-pool family inet]
user@host# set range lns-v4-pool-range low 192.168.1.1 high 192.168.255.255
```

NOTE: Dual-stack (IPv4/IPv6) LNS is supported, but IPv6-only subscribers are not supported.

To configure an IPv6 address-assignment pool for L2TP LNS:

1. Configure the name of the pool and specify the IPv6 family.

   [edit access]
   user@host# edit address-assignment pool *pool-name* family inet6

2. Configure the IPv6 network prefix for the address pool. The prefix specification is required when you configure an IPv6 address-assignment pool.

   [edit access address-assignment pool *pool-name* family inet6]
   user@host# set prefix *ipv6-prefix*

3. Configure the name of the range and define the range. You can define the range based on the lower and upper boundaries of the prefixes in the range, or based on the length of the prefixes in the range.

   [edit access address-assignment pool *pool-name* family inet6]
   user@host# set range *range-name* low *lower-limit* high *upper-limit*

For example, to configure an IPv6 address-assignment pool:

```
[edit access]
user@host# edit address-assignment pool lns-v6-pool family inet6
[edit access address-assignment pool lns-v6-pool family inet6]
user@host# set prefix 2010:9999::/32
[edit access address-assignment pool lns-v6-pool family inet6]
user@host# set range lns-v6-pool-range low 2010:9999:1::/48 high 2010:9999::ffff::/48
```

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- Address-Assignment Pools Overview on page 167
- Configuring Address-Assignment Pools on page 168

Configuring the L2TP LNS Peer Interface

The peer interface connects the LNS to the cloud towards the LACs so that IP packets can be exchanged between the tunnel endpoints. MPLS and aggregated Ethernet can also be used to reach the LACs.

NOTE: On MX Series routers, you must configure the peer interface on an MPC.

To configure the LNS peer interface:

1. Specify the interface name.

   [edit interfaces]
   user@host# edit *interface-name*

2. Enable VLANs.

   [edit interfaces *interface-name*]
   user@host# set vlan-tagging

3. Specify the logical interface, bind a VLAN tag ID to the interface, and configure the address family and the IP address for the logical interface.

   [edit interfaces *interface-name*]
   user@host# edit unit *logical-unit-number*
   [edit interfaces *interface-name* unit *logical-unit-number*]
   user@host# set vlan-id *number*
   user@host# set family *family* address *ip-address*

   NOTE: The IPv6 address family is not supported as a tunnel endpoint.

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Copyright © 2013, Juniper Networks, Inc.

Enabling Inline Service Interfaces

The inline service interface is a virtual physical interface that resides on the Packet Forwarding Engine. This **si** interface, referred to as an *anchor* interface, makes it possible to provide L2TP services without a special services PIC. The inline service interface is supported only by MPCs on MX Series routers. Four inline service interfaces are configurable per MPC-occupied chassis slot.

NOTE: On MX80 routers, you can configure only four inline services physical interfaces as anchor interfaces for L2TP LNS sessions: si-1/0/0, si-1/1/0, si-1/2/0, and si-1/3/0. You cannot configure si-0/0/0 for this purpose on MX80 routers.

To enable inline service interfaces:

1. Access an MPC-occupied slot and the PIC where the interface is to be enabled.

   [edit chassis]
   user@host# **edit fpc** ***slot-number*** **pic** ***number***

2. Enable the interface and specify the amount of bandwidth reserved on each Packet Forwarding Engine for tunnel traffic using inline services.

   [edit chassis fpc *slot-number* pic *number*]
   user@host# **set inline-services bandwidth (1g | 10g)**

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Configuring an Inline Service Interface for L2TP LNS

The inline service interface is a virtual physical service interface that resides on the Packet Forwarding Engine. This **si** interface, referred to as an *anchor* interface, makes it possible to provide L2TP services without a special services PIC. The inline service interface is supported only by MPCs on MX Series routers. Four inline service interfaces are configurable per MPC-occupied chassis slot.

You can maximize the number of sessions that can be shaped in one service interface by setting the maximum number of hierarchy levels to two. In this case, each LNS session consumes one L3 node in the scheduler hierarchy for shaping.

If you do not specify the number of levels (two is the only option), then the number of LNS sessions that can be shaped on the service interface is limited to the number of L2 nodes, or 4096 sessions. Additional sessions still come up, but they are not shaped.

To configure an inline service interface:

1. Access the service interface.

   [edit interfaces]
   user@host# **edit si-*slot/pic/port***

2. (Optional; for per-session shaping only) Enable the inline service interface for hierarchical schedulers and limit the number of scheduler levels to two.

   [edit interfaces si-*slot/pic/port*]
   user@host# **set hierarchical-scheduler maximum-hierarchy-levels 2**

3. (Optional; for per-session shaping only) Configure services encapsulation for inline service interface.

   [edit interfaces si-*slot/pic/port*]
   user@host# **set encapsulation generic-services**

4. Configure the IPv4 family on the reserved unit 0 logical interface.

   [edit interfaces si-*slot/pic/port*]
   user@host# **set unit 0 family inet**

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Configuring Options for the LNS Inline Services Logical Interface

You must specify characteristics—**dial-options**—for each of the inline services logical interfaces that you configure for the LNS. LNS on MX Series routers supports only one session per logical interface, so you must configure it as a **dedicated** interface; the **shared** option is not supported. (LNS on M Series routers supports **dedicated** and **shared** options.) You also configure an identifying name for the logical interface that matches the name you specify in the access profile.

To configure the logical interface options:

1. Access the inline services logical interface.

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

2. Specify an identifier for the logical interface.

   [edit interfaces si-*fpc/pic/port* unit *logical-unit-number*]
   user@host# **set dial-options l2tp-interface-id *name***

3. Configure the logical interface to be used for only one session at a time.

   [edit interfaces si-*fpc/pic/port* unit *logical-unit-number*]
   user@host# **set dial-options dedicated**

4. Configure the address family for each logical interface and enable the local address on the LNS that provides local termination for the L2TP tunnel to be derived from the specified interface name.

   [edit interfaces si-*fpc/pic/port* unit *logical-unit-number*]
   user@host# **set family inet unnumbered-address lo0.0**

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- Configuring an L2TP Access Profile on the LNS on page 422

Copyright © 2013, Juniper Networks, Inc.

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

The L2TP tunnel group specifies attributes that apply to L2TP tunnels and sessions from a group of LAC clients. These attributes include the access profile used to validate L2TP connection requests made to the LNS on the local gateway address, a local access profile that overrides the global access profile, the keepalive timer, and whether the IP ToS value is reflected.

NOTE: **If you delete a tunnel group, all L2TP sessions in that tunnel group are terminated. If you change the value of the local-gateway-address, service-device-pool, or service-interface statements, all L2TP sessions using those settings are terminated. If you change or delete other statements at the [edit services l2tp tunnel-group *name*] hierarchy level, new tunnels you establish use the updated values but existing tunnels and sessions are not affected.**

To configure the LNS tunnel group:

1. Create the tunnel group.

   [edit services l2tp]
   user@host# **edit** tunnel-group ***name***

2. Specify the service anchor interface responsible for L2TP processing on the LNS.

   [edit services l2tp tunnel-group *name*]
   user@host# **set** service-interface ***interface-name***

   This service anchor interface is required for static LNS sessions, and for dynamic LNS sessions that do not balance traffic across a pool of anchor interfaces. The interface is configured at the **[edit interfaces]** hierarchy level.

3. (Optional; for load-balancing dynamic LNS sessions only) Specify a pool of inline service anchor interfaces to enable load-balancing of L2TP traffic across the interfaces.

   [edit services l2tp tunnel-group *name*]
   user@host# **set** service-device-pool ***pool-name***

   The pool is defined at the **[edit services service-device-pools]** hierarchy level.

4. (For dynamic LNS sessions only) Specify the name of the dynamic profile that defines and instantiates inline service interfaces for L2TP tunnels

   [edit services l2tp tunnel-group *name*]
   user@host# **set** dynamic-profile ***profile-name***

   The profile is defined at the **[edit dynamic-profiles]** hierarchy level.

5. Specify the access profile that validates all L2TP connection requests to the local gateway address.

   [edit services l2tp tunnel-group *name*]
   user@host# **set** l2tp-access-profile ***profile-name***

Copyright © 2013, Juniper Networks, Inc.

6. Configure the local gateway address on the LNS; corresponds to the IP address that is used by LACs to identify the LNS.

   [edit services l2tp tunnel-group *name*]
   user@host# set local-gateway address *address*

7. (Optional) Configure the interval at which the LNS sends hello messages if it has received no messages from the LAC.

   [edit services l2tp tunnel-group *name*]
   user@host# set hello-interval *seconds*

8. (Optional) Specify a local access profile that overrides the global access profile to configure RADIUS server settings for the tunnel group.

   [edit services l2tp tunnel-group *name*]
   user@host# set aaa-access-profile *profile-name*

   This local profile is configured at the **[edit access profile]** hierarchy level.

9. (Optional) Configure the LNS to reflect the IP ToS value from the inner IP header to the outer IP header (applies to CoS configurations).

   [edit services l2tp tunnel-group *name*]
   user@host# set tos-reflect

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- Configuring an L2TP Access Profile on the LNS on page 422

Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions

You can create a pool of inline service interfaces, also known as a *service device pool*, to enable load-balancing of L2TP traffic across the interfaces. The pool is supported for dynamic LNS configurations, where it provides a set of logical interfaces that can be dynamically created and allocated to L2TP sessions on the LNS. The pool is assigned to an LNS tunnel group. L2TP maintains the state of each inline service interface and uses a round-robin method to evenly distribute the load among available interfaces when new session requests are accepted.

NOTE: **Load balancing is available only for dynamically created subscriber interfaces.**

LNS sessions anchored on an MPC are not affected by a MIC failure as long as some other path to the peer LACs exists. If the MPC hosting the peer interface fails and there is no path to peer LACs, the failure initiates termination and clean-up of all the sessions on the MPC.

If the MPC anchoring the LNS sessions itself fails, the Routing Engine does not relocate sessions to another slot and all sessions are terminated immediately. New sessions can come up on another available interface when the client retries.

 Copyright © 2013, Juniper Networks, Inc.

To configure the service device pool:

1. Create the pool.

   [edit services service-device-pools]
   user@host# edit pool *pool-name*

2. Specify the inline service interfaces that make up the pool.

   [edit services service-device-pools pool *pool-name*]
   user@host# set interface *service-interface-name*
   user@host# set interface *service-interface-name*

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429

Configuring a Dynamic Profile for Dynamic LNS Sessions

You can configure L2TP to dynamically assign inline service interfaces for L2TP tunnels. You must define one or more dynamic profiles and assign a profile to each tunnel group. Both IPv4-only and dual-stack IPv4/IPv6 interfaces are supported.

To configure the L2TP dynamic profile:

1. Create the dynamic profile.

   [edit]
   user@host# edit dynamic-profiles *profile-name*

2. Configure the interface to be dynamically assigned to the routing instance used by the tunneled PPP clients.

   [edit dynamic-profiles *profile-name* routing-instances “$junos-routing-instance”]
   user@host# set interface $junos-interface-name

3. Configure the routing options for access routes in the routing instance.

   [edit dynamic-profiles *profile-name* routing-instances “$junos-routing-instance” routing-options access]
   user@host# set route next-hop $junos-framed-route-nexthop
   user@host# set route metric $junos-framed-route-cost
   user@host# set route preference $junos-framed-route-distance

4. Configure the routing options for access-internal routes in the routing instance.

   [edit dynamic-profiles *profile-name* routing-instances “$junos-routing-instance” routing-options access-internal]
   user@host# set route $junos-subscriber-ip-address

5. Define the interfaces used by the dynamic profile. The variable is dynamically replaced by one of the configured inline service interfaces.

   [edit dynamic-profiles *profile-name*]
   user@host# set interfaces $junos-interface-ifd-name

6. Configure the inline services logical interfaces to be dynamically instantiated.

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

7. Specify an identifier for the logical interfaces.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit]
user@host# set dial-options l2tp-interface-id name
```

8. Configure each logical interface to be used for only one session at a time.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit]
user@host# set dial-options dedicated
```

9. Configure the address family for the logical interfaces and enable the local address on the LNS that provides local termination for the L2TP tunnel to be derived from the specified interface name.

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit]
user@host# set family inet unnumbered-address $junos-loopback-interface
```

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429

Configuring L2TP Tunnel Switching

L2TP tunnel switching enables a router configured as an LTS to forward PPP packets carried on one L2TP session to a second L2TP session terminated on a different LNS. To configure L2TP tunnel switching, you must define a tunnel switch profile and then assign that profile.

To define an L2TP tunnel switch profile:

1. Create the profile.

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

2. (Optional) Override the default actions taken for certain L2TP AVPs at the switching boundary.

```
[edit access tunnel-switch-profile profile-name]
user@host# set avp bearer-type action
user@host# set avp calling-number action
user@host# set avp cisco-nas-port-info action
```

3. Specify the tunnel profile that defines the tunnel to which the subscriber traffic is switched.

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

 Copyright © 2013, Juniper Networks, Inc.

4. (Optional) Apply the profile as a global default profile to switch packets from all incoming sessions from the LAC.

```
[edit services l2tp]
user@host1# set tunnel-switch-profile profile-name
```

5. (Optional) Apply the profile as part of a tunnel group to switch packets from all sessions in the tunnel group.

```
[edit services l2tp tunnel-group name]
user@host1# set tunnel-switch-profile profile-name
```

NOTE: The tunnel group is part of the LTS configuration that enables it to act as the LNS for the original sessions from the LAC.

A tunnel group with a tunnel switch profile must also contain a dynamic profile, because tunnel switching supports only dynamic subscribers.

6. (Optional) Apply the profile as part of a domain map to switch packets from all sessions that are associated with the domain.

```
[edit access domain map domain-map-name]
user@host1# set tunnel-switch-profile profile-name
```

NOTE: A domain map cannot have both a tunnel switch profile and a tunnel profile. You must remove one if you add the other.

7. (Optional) Apply the profile by means of the Tunnel-Switch-Profile VSA [26–91] in the RADIUS Access-Accept message returned when the session from the LAC is authenticated. Refer to the documentation for your RADIUS server to determine how to configure this method.

For example, consider the following configuration:

```
[edit access tunnel-switch-profile l2tp-tunnel-switch-profile]
user@host# set avp bearer-type regenerate
user@host# set avp calling-number regenerate
user@host# set avp cisco-nas-port-info drop
user@host# set tunnel-profile l2tp-tunnel-profile1

[edit access tunnel-switch-profile lts-profile-groupA]
user@host# set tunnel-profile l2tp-tunnel-profile2
[edit access tunnel-switch-profile lts-profile-example.com]
user@host# set tunnel-profile l2tp-tunnel-profile3

[edit services l2tp]
user@host1# set tunnel-switch-profile l2tp-tunnel-switch-profile
user@host1# set tunnel-group groupA tunnel-switch-profile lts-profile-groupA

[edit access domain]
user@host1# set map example.com tunnel-switch-profile lts-profile-example.com
```

This configuration creates three tunnel switch profiles, l2tp-tunnel-switch-profile, lts-profile-groupA, and lts-profile-example-com.

The profile l2tp-tunnel-switch-profile is applied as the global default. When packets are switched according to this profile, the values for the Bearer Type AVP (18) and Calling Number AVP (22) in the L2TP packets are regenerated based on local policy at the L2TP tunnel switch and then sent with the packets. The Cisco NAS Port Info AVP (100) is simply dropped. Finally, l2tp-tunnel-profile1 provides the configuration characteristics of the tunnel to which the traffic is switched.

Tunnel switch profile lts-profile-groupA is applied by means of a tunnel group, groupA; it specifies a different tunnel profile, l2tp-tunnel-profile2 and it does not override any AVP actions. Tunnel switch profile lts-profile-example.com is applied by means of a domain map for the example.com domain; it specifies a different tunnel profile, l2tp-tunnel-profile3 and it does not override any AVP actions.

Related Documentation

- L2TP Tunnel Switching Overview on page 392
- Tunnel Switching Actions for L2TP AVPs at the Switching Boundary on page 396

Tracing L2TP Operations for Subscriber Access

The Junos OS trace feature tracks L2TP operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

NOTE: **This topic refers to tracing L2TP operations on MX Series routers. To trace L2TP operations on M Series routers, see *Tracing L2TP Operations*.**

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file located in the **/var/log** directory. By default, the router uses the filename **jl2tpd**. You can specify a different filename, but you cannot change the directory in which trace files are located.
2. When the trace log file ***filename*** reaches 128 kilobytes (KB), it is compressed and renamed ***filename*.0.gz**. Subsequent events are logged in a new file called ***filename***, until it reaches capacity again. At this point, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until the number of archived files reaches the maximum file number. Then the oldest trace file—the one with the highest number—is overwritten.

   You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)

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 © 2013, Juniper Networks, Inc.

To configure L2TP tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the L2TP Trace Log Filename" on page 435.

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

   See "Configuring the Number and Size of L2TP Log Files" on page 435.

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

   See "Configuring Access to the L2TP Log File" on page 436.

4. (Optional) Configure a regular expression to filter the information to be included in the trace log.

   See "Configuring a Regular Expression for L2TP Messages to Be Logged" on page 436.

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

   See "Configuring the L2TP Tracing Flags" on page 437.

6. (Optional) Configure a severity level for messages to specify which event messages are logged.

   See "Configuring the Severity Level to Filter Which L2TP Messages Are Logged" on page 437.

Configuring the L2TP Trace Log Filename

By default, the name of the file that records trace output for L2TP is **jl2tpd**. You can specify a different name with the **file** option.

To configure the filename for L2TP tracing operations:

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

  ```
  [edit services l2tp traceoptions]
  user@host# set file l2tp_logfile_1
  ```

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 434

Configuring the Number and Size of L2TP Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format ***.number*.gz**. The newest archived file is **.0.gz** and the oldest archived file is ***.(maximum number)-1*.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output.

  [edit services l2tp traceoptions]
  user@host# **set file l2tp_1 _logfile_1 files 20 size 2097152**

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 434

Configuring Access to the L2TP Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

To specify that all users can read the log file:

- Configure the log file to be world-readable.

  [edit services l2tp traceoptions]
  user@host# **set file l2tp_1 _logfile_1 world-readable**

To explicitly set the default behavior, only the user who configured tracing can read the log file:

- Configure the log file to be no-world-readable.

  [edit services l2tp traceoptions]
  user@host# **set file l2tp_1 _logfile_1 no-world-readable**

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 434

Configuring a Regular Expression for L2TP Messages to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including regular expressions to be matched.

To configure regular expressions to be matched:

- Configure the regular expression.

  [edit services l2tp traceoptions]
  user@host# **set file l2tp_1 _logfile_1 match *regex***

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 434

Configuring the L2TP Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

  [edit services l2tp traceoptions]
  user@host# **set flag** ***flag***

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 434

Configuring the Severity Level to Filter Which L2TP Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. A low severity level is less restrictive—filters out fewer messages—than a higher level. When you configure a severity level, all messages at that level and all higher (more restrictive) levels are logged.

The following list presents severity levels in order from lowest (least restrictive) to highest (most restrictive). This order also represents the significance of the messages; for example, **error** messages are of greater concern than **info** messages.

- **verbose**
- **info**
- **notice**
- **warning**
- **error**

The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all**. You can also specify **verbose** with the same result, because **verbose** is the lowest (least restrictive) severity level; it has nothing to do with the terseness or verbosity of the messages. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages. By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

  [edit services l2tp traceoptions]

```
user@host# set level severity
```

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 434

Verifying and Managing L2TP for Subscriber Access

Purpose View or clear information about L2TP tunnels and sessions.

Action

- To display a summary of L2TP tunnels, sessions, errors, and control and data packets:

```
user@host> show services l2tp summary
```

- To display the L2TP destinations:

```
user@host> show services l2tp destination
```

- To clear all L2TP destinations:

```
user@host> clear services l2tp destination all
```

- To clear statistics for all L2TP tunnels belonging to a destination, tunnels belonging to a specified local-gateway address, and tunnels belonging to a specified peer-gateway address:

```
user@host>clear services l2tp destination statistics all
user@host>clear services l2tp destination local-gateway 10.1.1.2
```

- To display the L2TP sessions:

```
user@host> show services l2tp session
```

- To clear all L2TP sessions, the session with a specified local session ID, or sessions associated with the local gateway specified by an IP address or name:

```
user@host>clear services l2tp session all
user@host>clear services l2tp session local-session-id 40553
user@host>clear services l2tp session local-gateway 10.1.1.2
user@host>clear services l2tp session local-gateway-name lns-mx960
```

- To clear statistics for all L2TP sessions, the session with a specified local session ID, or sessions associated with the local gateway specified by an IP address or name:

```
user@host>clear services l2tp session statistics all
user@host>clear services l2tp session statistics local-session-id 17967
user@host>clear services l2tp session statistics local-gateway 10.1.1.2
user@host>clear services l2tp session statistics local-gateway-name lns-mx960
```

- To display the L2TP tunnels:

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

- To clear all L2TP tunnels, the tunnel with a specified local tunnel ID, or tunnels associated with the local gateway specified by an IP address or name:

```
user@host> clear services l2tp tunnel all
user@host>clear services l2tp tunnel local-tunnel-id 40553
user@host>clear services l2tp tunnel local-gateway 10.1.1.2
user@host>clear services l2tp tunnel local-gateway-name lns-mx960
```

 Copyright © 2013, Juniper Networks, Inc.

- To clear statistics for all L2TP tunnels, the tunnel with a specified local tunnel ID, or tunnels associated with the local gateway specified by an IP address or name:

  user@host> **clear services l2tp tunnel statistics all**
  user@host>**clear services l2tp tunnel statistics local-tunnel-id 40553**
  user@host>**clear services l2tp tunnel statistics local-gateway 10.1.1.2**
  user@host>**clear services l2tp tunnel statistics local-gateway-name lns-mx960**

Related Documentation

- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- *Junos OS Operational Mode Commands*.

Testing L2TP Tunnel Configurations from the LAC

You can test L2TP tunnel configurations on the LAC and successful subscriber authentication and tunneling without bringing up a PPP user and an associated tunnel.

Issue the **test services l2tp tunnel** command from CLI operational mode to map a subscriber to an L2TP tunnel, verify the L2TP tunnel configuration (both locally on the LAC and on a back-end server such as a RADIUS server), and verify that L2TP tunnels from the LAC can be established with the remote LNS.

The Junos OS LAC implementation enables you to configure multiple tunnels from which one tunnel is chosen for tunneling a PPP subscriber. You can use the **test services l2tp tunnel** command to test all possible tunnel configurations to verify that each can be established. Alternatively, you can test only a specific tunnel for the subscriber.

You must specify a configured subscriber username when you issue the command. The test generates a dummy password—*testpass*—for the subscriber, or you can optionally specify the password. The test verifies whether the subscriber identified by that username can be tunneled according to the tunnel configuration. If the subscriber can be tunneled, then the test verifies whether the L2TP tunnel can be established with the LNS according to the L2TP configuration.

You can optionally specify a tunnel ID, in which case only that tunnel is tested; the tunnel must be already configured for that username. If you omit this option, the test is applied to the full set of tunnel configurations that are returned for the username. The tunnel ID you specify is the same as that used by Tunnel-Assignment-Id (RADIUS attribute 82) and specified by the **identification** statement in the tunnel profile.

To test subscriber authentication and tunnel configuration:

- Specify only the username.

  Example 1:

  ```
  user@host> test services l2tp tunnel user test-user1@example.com
  Subscriber: test-user1@example.com, authentication failed
  ```

  The user failed authentication with the generated password and consequently was not tunneled.

Example 2:

```
user@host> test services l2tp tunnel user user23@example.com
Subscriber: user23@example.com, authentication success, l2tp tunneled
  Tunnel-name   Tunnel-peer   Logical-System  Routing-Instance Status
  test1tunnel    192.168.2.3      default         default         Up
  test2tunnel    172.24.3.3       default         default        Peer
unresponsive
  test3tunnel    172.24.5.1       default           test          Up
```

This user was authenticated with the generated password and successfully tunneled. A set of tunnels was found to be associated with that username and the entire set was tested.

- Specify the username and the user's configured password.

```
user@host> test services l2tp tunnel user test-user1@example.com password grZ98#jW
Subscriber: test-user1@example.com, authentication success, locally terminated
```

The subscriber was authenticated. However, the user was terminated locally rather than tunneled; this means that no tunnel was found to be associated with the user.

- Specify the username and a particular tunnel for the subscriber.

```
user@host> test services l2tp tunnel user rx37w@example.com tunnel-name ce-lac
Subscriber: rx37w@example.com, authentication success, l2tp tunneled
  Tunnel-name   Tunnel-peer   Logical-System  Routing-Instance Status
  ce-lac         192.168.5.10      default         default         Up
```

The subscriber was authenticated and tunneled. The specified tunnel was found for the subscriber and the tunnel was established, confirming the tunnel configuration.

- Specify the username, the user's configured password, and a tunnel.

```
user@host> test services l2tp tunnel user fanta4-mfg-fan@example.com password dieda499
tunnel-name tunnel5
Subscriber: fanta4-mfg-fan@example.com, authentication success, l2tp tunneled
```

The subscriber was authenticated and tunneled. The absence of tunnel information in the output indicates that the specified tunnel configuration does not exist.

Related Documentation

- L2TP for Subscriber Access Overview on page 379

Example: Configuring an L2TP LNS

This example shows how you can configure an L2TP LNS on an MX Series router to provide tunnel endpoints for an L2TP LAC in your network. This configuration includes a dynamic profile for dual-stack subscribers.

- Requirements on page 441
- Overview on page 442
- Configuration on page 443

 Copyright © 2013, Juniper Networks, Inc.

Requirements

L2TP LNS requires the following hardware and software:

- MX Series 3D Universal Edge Router
- One or more MPCs
- Junos OS Release 11.4 or later

No special configuration beyond device initialization is required before you can configure this feature.

You must configure certain standard RADIUS attributes and Juniper Networks VSAs in the attribute return list on the AAA server associated with the LNS for this example to work. Table 46 on page 441 lists the attributes with their required order setting and values. We recommend that you use the most current Juniper Networks RADIUS dictionary, available in the *Downloads* box on the *Junos OS Subscriber Management* page for the current release at http://www.juniper.net/techpubs/en_US/junos/information-products/pathway-pages/subscriber-access/index.html.

Table 46: VSA and Standard RADIUS Attribute Names, Order, and Values Required for Example

VSA Name [Number]	Order	Value
CoS-Parameter-Type [26–108]	1	T01 Multiplay
CoS-Parameter-Type [26–108]	2	T02 10m
CoS-Parameter-Type [26–108]	3	T08 -36
CoS-Parameter-Type [26–108]	4	T07 cell-mode
Framed-IPv6-Pool [100]	0	jnpr_ipv6_pool
Framed-Pool [88]	0	jnpr_pool
Egress-Policy-Name [26-11]	0	classify
Ingress-Policy-Name [26-10]	0	classify
Virtual-Router [26-1]	0	default

Copyright © 2013, Juniper Networks, Inc.

Overview

The LNS employs user group profiles to apply PPP attributes to the PPP subscribers that are tunneled from the LAC. LACs in the network are clients of the LNS. The clients are associated with user group profiles in the L2TP access profile configured on the LNS. In this example, the user group profile **ce-l2tp-group-profile** specifies the following PPP attributes:

- A 30-second interval between PPP keepalive messages for L2TP tunnels from the client LAC terminating on the LNS.
- A 200-second interval that defines how long the PPP subscriber session can be idle before it is considered to have timed out.
- Both PAP and CHAP as the PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

The L2TP access profile **ce-l2tp-profile** defines a set of L2TP parameters for each client LAC. In this example, the user group profile **ce-l2tp-group-profile** is associated with both clients, **lac1** and **lac2**. Both clients are configured to have the LNS renegotiate the link control protocol (LCP) with the PPP client rather than accepting the pre-negotiated LCP parameters that the LACs pass to the LNS. LCP renegotiation also causes authentication to be renegotiated by the LNS; the authentication method is specified in the user group profile. The maximum number of sessions allowed per tunnel is set to 1000 for **lac1** and to 4000 for **lac2**. A different password is configured for each LAC.

A local AAA access profile, **aaa-profile**, enables you to override the global AAA access profile, so that you can specify an authentication order, a RADIUS server that you want to use for L2TP, and a password for the server.

In this example, an address pool defines a range of IP addresses that the LNS allocates to the tunneled PPP sessions. This example defines ranges of IPv4 and IPv6 addresses.

Two inline service interfaces are enabled on the MPC located in slot 5 of the router. For each interface, 10 Gbps of bandwidth is reserved for tunnel traffic on the interface's associated PFE. These *anchor* interfaces serve as the underlying physical interface. To enable CoS queue support on the individual logical inline service interfaces, you must configure both services encapsulation (**generic-services**) and hierarchical scheduling support on the anchors. The IPv4 address family is configured for both anchor interfaces. Both anchor interfaces are specified in the **lns_p1** service device pool. The LNS can balance traffic loads across the two anchor interfaces when the tunnel group includes the pool.

This example uses the dynamic profile **dyn-lns-profile2** to specify characteristics of the L2TP sessions that are created or assigned dynamically when a subscriber is tunneled to the LNS. For many of the characteristics, a predefined variable is set; the variables are dynamically replaced with the appropriate values when a subscriber is tunneled to the LNS.

The interface to which the tunneled PPP client connects (**$junos-interface-name**) is dynamically created in the routing instance (**$junos-routing-instance**) assigned to the subscriber. Routing options for access routes include the route's next hop address

 Copyright © 2013, Juniper Networks, Inc.

(**$junos-framed-route-nexthop**), metric (**$junos-framed-route-cost**), and preference (**$junos-framed-route-distance**). For access-internal routes, a dynamic IP address variable (**$junos-subscriber-ip-address**) is set.

The logical inline service interfaces are defined by the name of a configured anchor interface (**$junos-interface-ifd-name**) and a logical unit number (**$junos-interface-unit**). The profile assigns **l2tp-encapuslation** as the identifier for the logical interface and specifies that each interface can be used for only a single session at a time.

The IPv4 address is set to a value returned from the AAA server. For IPv4 traffic an input firewall filter **$junos-input-filter** and an output firewall filter **$junos-output-filter** are attached to the interface. The loopback variable (**$junos-loopback-interface**) derives an IP address from a loopback interface (**lo**) configured in the routing instance and uses it in IPCP negotiation as the PPP server address. Because this is a dual-stack configuration, the IPv6 address family is also set, with the addresses provided by the **$junos-ipv6-address** variable.

The **$junos-ipv6-address** variable is used because Router Advertisement Protocol is also configured. This variable enables AAA to allocate the first address in the prefix to be reserved as the local address for the interface. The minimal configuration for the Router Advertisement Protocol in the dynamic profile specifies the **$junos-interface-name** and **$junos-ipv6-ndra-prefix** variables to dynamically assign a prefix value in IPv6 neighbor discovery router advertisements.

The dynamic profile also includes the class of service configuration that is applied to the tunnel traffic. The traffic control profile (**tc-profile**) includes variables for the scheduler map (**$junos-cos-scheduler-map**), shaping rate (**$junos-cos-shaping-rate**), overhead accounting (**$junos-cos-shaping-mode**), and byte adjustment **$junos-cos-byte-adjust**). The dynamic profile applies the CoS configuration—including the forwarding class, the output traffic control profile, and the rewrite rules—to the dynamic service interfaces.

The **tg-dynamic** tunnel group configuration specifies the access profile **ce-l2tp-profile**, the local AAA profile **aaa-profile**, and the dynamic profile **dyn-lns-profile2** that are used to dynamically create LNS sessions and define the characteristics of the sessions. The **lns_p1** service device pool associates a pool of service interfaces with the group to enable LNS to balance traffic across the interfaces. The local gateway address **11.1.1.2** corresponds to the remote gateway address that is configured on the LAC.

NOTE: **This example does not show all possible configuration choices.**

Configuration

CLI Quick Configuration

To quickly configure an L2TP LNS, copy the following commands, paste them in a text file, remove any line breaks, and then copy and paste the commands into the CLI.

```
[edit]
edit access group-profile ce-l2tp-group-profile
set ppp idle-timeout 200
set ppp ppp-options pap
set ppp ppp-options chap
```

```
set ppp keepalive 30
top
edit access profile ce-l2tp-profile
set client lac1 l2tp maximum-sessions-per-tunnel 1000
set client lac1 l2tp interface-id l2tp-encapsulation-1
set client lac1 l2tp lcp-renegotiation
set client lac1 l2tp shared-secret "lac1-secret"
set client lac1 user-group-profile ce-l2tp-group-profile
set client lac2 l2tp maximum-sessions-per-tunnel 4000
set client lac2 l2tp interface-id l2tp-encap-2
set client lac2 l2tp lcp-renegotiation
set client lac2 l2tp shared-secret "lac2-secret"
set client lac2 user-group-profile ce-l2tp-group-profile
top
edit access profile aaa-profile
set authentication-order radius
set radius-server 172.21.146.93 secret "aaa-secret”
top
edit access address-assignment pool client-pool1 family inet
set network 192.168.1.1/16
set range lns-v4-pool-range low 192.168.1.1
set range lns-v4-pool-range high 192.168.255.255
top
edit access address-assignment pool client-ipv6-pool2 family inet6
set prefix 2010:db8::/32
set range lns-v6-pool-range low 2010:db8:1::/48
set range lns-v6-pool-range high 2010:db8:ffff::/48
top
set interfaces ge-5/0/1 unit 11 vlan-id 11
set interfaces ge-5/0/1 unit 11 family inet address 11.1.1.2/24
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
top
set chassis fpc 5 pic 0 inline-services bandwidth 10g
set chassis fpc 5 pic 2 inline-services bandwidth 10g
top
edit interfaces si-5/0/0
set hierarchical-scheduler maximum-hierarchy-levels 2
set encapsulation generic-services
set unit 0 family inet
top
edit interfaces si-5/2/0
set hierarchical-scheduler maximum-hierarchy-levels 2
set encapsulation generic-services
set unit 0 family inet
top
set services service-device-pools pool lns_p1 interface si-5/0/0
set services service-device-pools pool lns_p1 interface si-5/2/0
top
edit dynamic-profiles dyn-lns-profile2 routing-instances $junos-routing-instance
set interface $junos-interface-name
edit routing-options access route $junos-framed-route-ip-address-prefix
set next-hop $junos-framed-route-nexthop
set metric $junos-framed-route-cost
set preference $junos-framed-route-distance
up 2
edit access-internal route $junos-subscriber-ip-address
```

 Copyright © 2013, Juniper Networks, Inc.

```
set qualified-next-hop $junos-interface-name
up 5
edit interfaces $junos-interface-ifd-name unit $junos-interface-unit
set dial-options l2tp-interface-id l2tp-encapsulation
set dial-options dedicated
set family inet filter input $junos-input-filter
set family inet filter output $junos-output-filter
set family inet unnumbered-address $junos-loopback-interface
set family inet6 address $junos-ipv6-address
set family inet6 filter input $junos-input-ipv6-filter
set family inet6 filter output $junos-output-ipv6-filter
up 3
edit protocols router-advertisement
set interface $junos-interface-name prefix $junos-ipv6-ndra-prefix
top
[edit class-of-service]
edit rewrite-rules dscp rewriteDSCP forwarding-class expedited-forwarding
set loss-priority high code-point af11
set loss-priority high code-point af12
top
edit dynamic-profiles dyn-lns-profile2 class-of-service traffic-control-profiles tc-profile
set scheduler-map $junos-cos-scheduler-map
set shaping-rate $junos-cos-shaping-rate
set overhead-accounting $junos-cos-shaping-mode
set overhead-accounting bytes $junos-cos-byte-adjust
up
edit interfaces $junos-interface-ifd-name unit $junos-interface-unit
set forwarding-class expedited-forwarding
set output-traffic-control-profile tc-profile
set rewrite-rules dscp rewriteDSCP
edit interfaces si-5/0/0
set output-control-profile-remaining tc-profile
top
set services l2tp tunnel-group tg-dynamic l2tp-access-profile ce-l2tp-profile
set services l2tp tunnel-group tg-dynamic aaa-access-profile aaa-profile
set services l2tp tunnel-group tg-dynamic local-gateway address 11.1.1.2
set services l2tp tunnel-group tg-dynamic service-device-pool lns_p1
set services l2tp tunnel-group tg-dynamic dynamic-profile dyn-lns-profile2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an L2TP LNS with inline service interfaces:

1. Configure a user group profile that defines the PPP configuration for tunnel subscribers.

```
[edit access]
user@host# edit group-profile ce-l2tp-group-profile
[edit access group-profile ce-l2tp-group-profile]
user@host# set ppp keepalive 30
user@host# set ppp idle-timeout 200
user@host# set ppp ppp-options chap
user@host# set ppp ppp-options pap
```

2. Configure an L2TP access profile that defines the L2TP parameters for each client LAC. This includes associating a user group profile with the client and specifying the identifier for the inline services logical interface that represents an L2TP session on the LNS.

```
[edit access profile ce-l2tp-profile client lac1]
user@host# set l2tp interface-id l2tp-encapsulation
user@host# set l2tp maximum-sessions-per-tunnel 1000
user@host# set l2tp shared-secret "lac1-secret"
user@host# set l2tp lcp-renegotiation
user@host# set user-group-profile ce-l2tp-group-profile
[edit access profile ce-l2tp-profile client lac2]
user@host# set l2tp interface-id interface-id
user@host# set l2tp maximum-sessions-per-tunnel 4000
user@host# set l2tp shared-secret "lac2-secret"
user@host# set l2tp lcp-renegotiation
user@host# set user-group-profile ce-l2tp-group-profile
```

3. Configure a AAA access profile to override the global access profile for the order of AAA authentication methods and server attributes.

```
[edit access profile aaa-profile]
user@host# set authentication-order radius
user@host# set radius-server 172.21.146.93 secret "aaa-secret”
```

4. Configure IPv4 and IPv6 address-assignment pools to allocate addresses for the clients (LACs).

```
[edit access address-assignment pool client-pool1 family inet]
user@host# set network 192.168.1.1/16
user@host# set range lns-v4-pool-range low 192.168.1.1 high 192.168.255.255
[edit access address-assignment pool client-ipv6-pool2 family inet6]
user@host# set prefix 2010:DB8::/32
user@host# set range lns-v6-pool-range low 2010:DB8:1::/48
user@host# set range lns-v6-pool-range high 2010:DB8:ffff::/48
```

5. Configure the peer interface to terminate the tunnel and the PPP server-side IPCP address (loopback address).

```
[edit interfaces ge-5/0/1
user@host# set vlan-tagging
user@host# set unit 11
[edit interfaces ge-5/0/1.11
user@host# set vlan-id 11
user@host# set family inet address 11.1.1.2/24
[edit interfaces lo0]
user@host# set unit 0 family inet address 127.0.0.1/32
```

6. Enable inline service interfaces on an MPC.

```
[edit chassis fpc 5]
user@host# set pic 0 inline-services bandwidth 10g
user@host# set pic 2 inline-services bandwidth 10g
```

7. Configure the anchor service interfaces with services encapsulation, hierarchical scheduling, and the address family.

```
[edit interfaces si-5/0/0]
user@host# set hierarchical-scheduler maximum hierarchy-levels 2
```

 Copyright © 2013, Juniper Networks, Inc.

```
user@host# set encapsulation generic-services
user@host# set unit 0 family inet
[edit interfaces si-5/2/0]
user@host# set hierarchical-scheduler maximum hierarchy-levels 2
user@host# set encapsulation generic-services
user@host# set unit 0 family inet
```

8. Configure a pool of service interfaces for dynamic LNS sessions.

```
[edit services service-device-pools pool lns_p1]
user@host# set interface si-5/0/0
user@host# set interface si-5/2/0
```

9. Configure a dynamic profile that dynamically creates L2TP logical interfaces for dual-stack subscribers.

```
[edit dynamic-profiles dyn-lns-profile2]
user@host# edit routing-instances $junos-routing-instance
user@host# set interface $junos-interface-name
[edit dynamic-profiles dyn-lns-profile2 routing-instances "$junos-routing-instance"]
user@host# edit routing-options access route $junos-framed-route-ip-address-prefix
[edit dynamic-profiles dyn-lns-profile2 routing-instances "$junos-routing-instance"
  routing-options access route "$junos-framed-route-ip-address-prefix"]
user@host# set next-hop $junos-framed-route-nexthop
user@host# set metric $junos-framed-route-cost
user@host# set preference $junos-framed-route-distance
[edit dynamic-profiles dyn-lns-profile2 routing-instances "$junos-routing-instance"
  routing-options access-internal]
user@host# set route $junos-subscriber-ip-address qualified-next-hop
  $junos-interface-name
[edit dynamic-profiles dyn-lns-profile2 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set dial-options l2tp-interface-id l2tp-encapsulation
user@host# set dial-options dedicated
user@host# set family inet unnumbered-address $junos-loopback-interface
user@host# set family inet filter input $junos-input-filter
user@host# set family inet filter output $junos-output-filter
user@host# set family inet6 address $junos-ipv6-address
set family inet6 filter input $junos-input-ipv6-filter
set family inet6 filter output $junos-output-ipv6-filter
[edit dynamic-profiles dyn-lns-profile2 protocols router-advertisement]
user@host# set interface $junos-interface-name prefix $junos-ipv6-ndra-prefix
```

10. Configure shaping, scheduling, and rewrite rules, and apply in the dynamic profile to tunnel traffic.

```
[edit class-of-service]
user@host# edit rewrite-rules dscp rewriteDSCP forwarding-class
  expedited-forwarding
user@host# set loss-priority high code-point af11
user@host# set loss-priority high code-point af12
[edit dynamic-profiles dyn-lns-profile2 class-of-service traffic-control-profiles
  tc-profile]
user@host# set scheduler-map $junos-cos-scheduler-map
user@host# set shaping-rate $junos-cos-shaping-rate
user@host# set overhead-accounting $junos-cos-shaping-mode
user@host# set overhead-accounting bytes $junos-cos-byte-adjust
```

```
[edit dynamic-profiles dyn-lns-profile2 class-of-service interfaces
  “$junos-interface-ifd-name” unit "$junos-interface-unit"]
user@host# set forwarding-class expedited-forwarding
user@host# set output-traffic-control-profile tc-profile
user@host# set rewrite-rules dscp rewriteDSCP
[edit class-of-service interfaces si-5/0/0]
user@host# set output-traffic-control-profile-remaining tc-profile
```

11. Configure the L2TP tunnel group to bring up dynamic LNS sessions using the pool of inline service interfaces to enable load-balancing.

```
[edit services l2tp tunnel-group tg-dynamic]
user@host# set l2tp-access-profile ce-l2tp-profile
user@host# set local-gateway address 11.1.1.2
user@host# set aaa-access-profile aaa-profile
user@host# set dynamic-profile dyn-lns-profile2
user@host# set service-device-pool lns_p1
```

Results From configuration mode, confirm the access profile, group profile, AAA profile, and address-assignment pools configuration by entering the **show access** command. Confirm the inline services configuration by entering the **show chassis** command. Confirm the interface configuration by entering the **show interfaces** command. Confirm the dynamic profile configuration by entering the **show dynamic-profiles** command. Confirm the tunnel group configuration by entering the **show services l2tp** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show access
group-profile ce-l2tp-group-profile {
  ppp {
    idle-timeout 200;
    ppp-options {
      pap;
      chap;
    }
    keepalive 30;
  }
}
profile ce-l2tp-profile {
  client lac1 {
    l2tp {
      maximum-sessions-per-tunnel 1000;
      interface-id l2tp-encapsulation-1;
      lcp-renegotiation;
      shared-secret "$9$ZJGi.Pfz6/tmPtu1IleLxNbwgaZjmPQDi"; ## SECRET-DATA
    }
    user-group-profile ce-l2tp-group-profile;
  }
  client lac2 {
    l2tp {
      maximum-sessions-per-tunnel 4000;
      interface-id l2tp-encap-2;
      lcp-renegotiation;
      shared-secret ""$9$KCjvLNdVYoaUdVDi.m3ntuOREyevLdVY8X"; ## SECRET-DATA
```

 Copyright © 2013, Juniper Networks, Inc.

```
        }
        user-group-profile ce-l2tp-group-profile;
    }
}
profile aaa-profile {
    authentication-order radius;
    radius-server {
        172.21.146.93 secret "$9$41JZjk.5Qz6k."; ## SECRET-DATA
    }
}
address-assignment {
    pool client-pool1 {
        family inet {
            network 192.168.1.1/16;
            range lns-v4-pool-range {
                low 192.168.1.1;
                high 192.168.255.255;
            }
        }
    }
    pool client-ipv6-pool2 {
        family inet6 {
            prefix 2010:db8::/32;
            range lns-v6-pool-range {
                low 2010:db8:1::/48;
                high 2010:db8:ffff::/48;
            }
        }
    }
}

[edit]
user@host# show chassis
fpc 5 {
    pic 0 {
        inline-services {
            bandwidth 10g;
        }
    }
    pic 2 {
        inline-services {
            bandwidth 10g;
        }
    }
}

[edit]
user@host# show interfaces
ge-5/0/1 {
    vlan-tagging;;
    unit 11 {
        vlan-id 11;
        family inet {
            address 11.1.1.2/24;
        }
```

```
    }
  }
  si-5/0/0 {
    hierarchical-scheduler maximum-hierarchy-levels 2;
    encapsulation generic-services;
    unit 0 {
      family inet;
    }
  }
  si-5/2/0 {
    hierarchical-scheduler maximum-hierarchy-levels 2;
    encapsulation generic-services;
    unit 0 {
      family inet;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 127.0.0.1/32;
      }
    }
  }
  [edit]
  user@host# show dynamic-profiles
  dyn-lns-profile2 {
    routing-instances {
      "$junos-routing-instance" {
        interface "$junos-interface-name";
        routing-options {
          access {
            route $junos-framed-route-ip-address-prefix {
              next-hop "$junos-framed-route-nexthop";
              metric "$junos-framed-route-cost";
              preference "$junos-framed-route-distance";
            }
          }
          access-internal {
            route $junos-subscriber-ip-address {
              qualified-next-hop "$junos-interface-name";
            }
          }
        }
      }
    }
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
          dial-options {
            l2tp-interface-id l2tp-encapsulation;
            dedicated;
          }
          family inet {
            filter {
              input "$junos-input-filter";
              output "$junos-output-filter";
```

 Copyright © 2013, Juniper Networks, Inc.

```
                }
                unnumbered-address "$junos-loopback-interface";
            }
            family inet6 {
                address $junos-ipv6-address;
                input $junos-input-ipv6-filter;
                output $junos-output-ipv6-filter;
            }
        }
    }
}
protocols {
    router-advertisement {
        interface "$junos-interface-name" {
            prefix $junos-ipv6-ndra-prefix;
        }
    }
}
class-of-service {
    rewrite-rules {
        dscp rewriteDSCP {
            forwarding-class expedited-forwarding {
                loss-priority high code-point af11
                loss-priority high code-point af12
            }
        }
    }
    traffic-control-profiles {
        tc-profile {
            scheduler-map "$junos-cos-scheduler-map";
            shaping-rate "$junos-cos-shaping-rate";
            overhead-accounting "$junos-cos-shaping-mode" bytes "$junos-cos-byte-adjust";
        }
    }
    interfaces {
        "$junos-interface-ifd-name" {
            unit "$junos-interface-unit" {
                forwarding-class expedited-forwarding;
                output-traffic-control-profile tc-profile;
                rewrite-rules {
                    dscp rewriteDSCP;
                }
            }
        }
    }
}
}

[edit]
user@host# show services l2tp
tunnel-group tg-dynamic {
    l2tp-access-profile ce-l2tp-profile;
    aaa-access-profile aaa-profile;
    local-gateway {
        address 11.1.1.2;
```

```
    }
    service-device-pool lns_p1;
    dynamic-profile dyn-lns-profile2;
  }
```

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

Related Documentation

- L2TP for Subscriber Access Overview on page 379
- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- Configuring an L2TP LAC on page 402

Copyright © 2013, Juniper Networks, Inc.

PART 7

Diameter Base Protocol and Applications for Subscriber Access

- Diameter Base Protocol Overview on page 455
- Configuring Diameter Base Protocol on page 471
- JSRC and Juniper Networks Session Resource Control (SRC) Overview on page 485
- Configuring JSRC for Subscriber Access on page 491
- Subscribers on Static Interfaces on page 495
- Configuring Subscribers over Static Interfaces on page 499
- PTSP and Juniper Networks Session and Resource Control (SRC) on page 515
- Configuring the PTSP Application on page 525
- Configuring Packet-Triggered Subscriber Services on page 529
- Gx-Plus Overview on page 541
- Configuring Gx-Plus on page 549

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 17

Diameter Base Protocol Overview

- Diameter Base Protocol Overview on page 455
- Messages Used by Diameter Applications on page 457
- Diameter AVPs and Diameter Applications on page 462

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, and PTSP.

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.

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.

Copyright © 2013, Juniper Networks, Inc.

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.

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. Include the **route** statement at the **[edit diameter network-element *element-name* forwarding]** hierarchy level.

To configure a route for the DNE, include the **destination** (optional), **function** (optional), and **metric** statements at the **[edit diameter network-element *element-name* forwarding route *dne-route-name*]** hierarchy level.

Specify the Diameter peers associated with the DNE by including one or more **peer** statements at the **[edit diameter network-element *element-name*]** hierarchy level.

Set the priority for each peer with the **priority** statement at the **[edit diameter network-element *element-name* peer *peer-name*]** hierarchy level.

Diameter requires you to configure information about the origin node; this is the endpoint node that originates Diameter for the Diameter instance. Include the **host** and **realm** statements at the **[edit diameter]** hierarchy level to configure the Diameter origin.

You can optionally configure one or more *transports* to specify the source (local) address of the transport layer connection. To configure a Diameter transport, include the **transport** statement at the **[edit diameter]** hierarchy level. Then include the **address** statement at the **[edit diameter transport *transport-name*]** hierarchy level.

You can optionally specify a logical system and routing instance for the connection by including the **logical-system** and **routing-instance** statements at the **[edit diameter transport *transport-name*]** hierarchy level. By default, Diameter uses the *default* logical system and *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.

 Copyright © 2013, Juniper Networks, Inc.

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.

Related Documentation

- Configuring Diameter on page 471
- Messages Used by Diameter Applications on page 457
- Diameter AVPs and Diameter Applications on page 462
- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 485
- Juniper Networks Session and Resource Control (SRC) and PTSP Overview on page 516
- Gx-Plus for Provisioning Subscribers Overview on page 541

Messages Used by Diameter Applications

The following Diameter applications are supported by Junos OS:

- JSRC—A Juniper Networks Diameter application registered with the IANA (http://www.iana.org) as Juniper Policy-Control-JSRC, with an ID of 16777244. 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. Communicates with the SAE (remote SRC peer).
- 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). 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 decision, 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.

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 47 on page 457 describes Diameter messages the applications use.

Table 47: Diameter Messages and Diameter Applications

Diameter Message	Code	Application	Description
AA-Request (AAR)	265	JSRC, 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.

Table 47: Diameter Messages and Diameter Applications *(continued)*

Diameter Message	Code	Application	Description
AA-Answer (AAA)	265	JSRC, PTSP	Response from the SAE to the application's AA-Request message.
Abort-Session-Request (ASR)	274	JSRC, PTSP	Request from the SAE to the application to log out a provisioned subscriber.
Abort-Session-Answer (ASA)	274	JSRC, 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 © 2013, Juniper Networks, Inc.

Table 47: 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 © 2013, Juniper Networks, Inc.

Table 47: 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 © 2013, Juniper Networks, Inc.

Table 47: 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, PTSP	Notification from the router to the SAE that a provisioned subscriber has logged out.
Session-Termination-Answer (STA)	275	JSRC, PTSP	Response from the SAE to the router's STR message. Includes success or failure notification.

Related Documentation

- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 485
- Understanding JSRC-SAE Interactions on page 488
- Juniper Networks Session and Resource Control (SRC) and PTSP Overview on page 516
- Understanding PTSP-SAE Interactions on page 517
- Gx-Plus for Provisioning Subscribers Overview on page 541
- Understanding Gx-Plus Interactions Between the Router and the PCRF on page 543

Diameter AVPs and Diameter Applications

Diameter conveys information by including various attribute-value pairs (AVPs) in Diameter messages. Table 48 on page 462 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.

Table 48: Standard Diameter AVPs

Attribute Number	Diameter AVP	Application	Description	Type
1	User-Name	Gx-Plus, JSRC	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
8	Framed-IP-Address	Gx-Plus, JSRC, PTSP	Identifies the IPv4 address configured for the subscriber. This is the same value as for RADIUS Framed-IP-Address attribute [8].	OctetString
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
85	Acct-Interim-Interval	JSRC, PTSP	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, 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
263	Session-ID	Gx-Plus, JSRC, PTSP	Specifies the subscriber session identifier. The router assigns the value to uniquely identify a subscriber session.	UTF8String

Copyright © 2013, Juniper Networks, Inc.

Table 48: Standard Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
268	Result-Code	Gx-Plus, JSRC, 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

Table 48: Standard Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
277	Auth-Session-State	JSRC, PTSP	Indicates whether AAA session state is maintained. • 0—STATE MAINTAINED • 1—NO STATE MAINTAINED	Enumerated
295	Termination-Cause	JSRC, 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
415	CC-Request-Number	Gx-Plus	Identifies a request within a session. The combination of Session-Id and CC-Request-Type is globally unique. The number is incremented for each request during the course of a session. The number is reset when a router high availability event takes place.	Unsigned32
416	CC-Request-Type	Gx-Plus	Specifies the type of credit control request: • INITIAL REQUEST (1) • UPDATE REQUEST (2) • TERMINATION_REQUEST (3) • EVENT REQUEST (4)	Enumerated
431	Granted-Service-Unit	Gx-Plus	Contains the amount that can be provided of one or more of the following requested units specified by the client: CC-Input-Octets, CC-Output-Octets, CC-Time, or CC-Total-Octets. Included in CCA-I messages, and may be included in CCA-U messages.	Grouped
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

Copyright © 2013, Juniper Networks, Inc.

Table 48: Standard Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
480	Accounting-Record-Type	JSRC, PTSP	Specifies the type of account record for service accounting: • INTERIM_RECORD—Accounting record sent between the start and stop records, at intervals specified by the Acct-Interim-Interval AVP (AVP code 85). It contains cumulative accounting data for the existing accounting session. • START_RECORD—Accounting record sent when the service is activated to initiate the accounting session. It contains accounting data relevant to the initiation of that session. • STOP_RECORD—Accounting record sent when the service is deactivated to terminate the accounting session. It contains cumulative data relevant to that session.	Enumerated
1001	Charging-Rule-Install	Gx-Plus	Requests the installation of the rule (activation of the service) designated by the included Charging-Rule-Name AVP (1005). This AVP has a vendor ID of 10415 (3GPP).	Grouped
1002	Charging-Rule-Remove	Gx-Plus	Requests the removal of the rule (deactivation of the service) designated by the included Charging-Rule-Name AVP (1005). This AVP has a vendor ID of 10415 (3GPP).	Grouped
1005	Charging-Rule-Name	Gx-Plus	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

Copyright © 2013, Juniper Networks, Inc.

Juniper Networks AVPs are used in addition to the standard Diameter AVPs. These AVPs have an enterprise number of 2636. Table 49 on page 466 lists the Juniper Networks AVPs that the supported Diameter applications use.

Table 49: 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
2023	Juniper-Template-Name	JSRC, PTSP	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

Copyright © 2013, Juniper Networks, Inc.

Table 49: Juniper Networks Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
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
2053	Juniper-Acct-Record	JSRC, PTSP	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

Table 49: Juniper Networks Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 49: Juniper Networks Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
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
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 50 on page 469 lists the Tekelec AVPs. These four variables are used to provide substitution values for user-defined CoS service variables.

Table 50: 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

Copyright © 2013, Juniper Networks, Inc.

Table 50: Tekelec Diameter AVPs *(continued)*

Attribute Number	Diameter AVP	Application	Description	Type
5557	Tekelec-Charging-Rule-Argument	Gx-Plus	Defines the substitution attributes used to replace service variables. Includes Tekelec-Charging-Rule-Argument-Name AVP (5555) and Tekelec-Charging-Rule-Argument-Value AVP (5556).	Grouped
5558	Tekelec-Charging-Rule-With-Arguments	Gx-Plus	Requests the installation of the rule (activation of the service) designated by the included Charging-Rule-Name AVP (1005). Requested service variable substitutions are provided by the optionally included Tekelec-Charging-Rule-Argument AVP (5557).	Grouped

Related Documentation

- Understanding JSRC-SAE Interactions on page 488
- Understanding PTSP-SAE Interactions on page 517
- Understanding Gx-Plus Interactions Between the Router and the PCRF on page 543
- Diameter Base Protocol Overview on page 455
- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 485
- Juniper Networks Session and Resource Control (SRC) and PTSP Overview on page 516
- Gx-Plus for Provisioning Subscribers Overview on page 541

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 18

Configuring Diameter Base Protocol

- Configuring Diameter on page 471
- Configuring the Origin Attributes of the Diameter Instance on page 472
- Configuring Diameter Peers on page 473
- Configuring Diameter Network Elements on page 473
- Configuring the Diameter Transport on page 474
- Tracing Diameter Base Protocol Processes for Subscriber Access on page 475
- Configuring the Diameter Base Protocol Trace Log Filename on page 476
- Configuring the Number and Size of Diameter Base Protocol Log Files on page 476
- Configuring Access to the Diameter Base Protocol Log File on page 477
- Configuring a Regular Expression for Diameter Base Protocol Messages to Be Logged on page 477
- Configuring the Diameter Base Protocol Tracing Flags on page 478
- Configuring the Severity Level to Filter Which Diameter Base Protocol Messages Are Logged on page 478
- Verifying Diameter Node, Instance, and Route Information on page 478
- Verifying and Managing Diameter Function Information on page 479
- Verifying and Managing Diameter Peer Information on page 480
- Verifying Diameter Network Element Information on page 481
- Troubleshooting Diameter Network Configuration on page 482
- Troubleshooting Diameter Network Connectivity on page 482

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 472

2. Configure the Diameter peers.

   See "Configuring Diameter Peers" on page 473

3. (Optional) Configure the Diameter transport layer elements.

   See "Configuring the Diameter Transport" on page 474

4. (Optional) Configure the Diameter network elements.

   See "Configuring Diameter Network Elements" on page 473

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

   See "Tracing Diameter Base Protocol Processes for Subscriber Access" on page 475.

Related Documentation

- Diameter Base Protocol Overview on page 455

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 471
- origin on page 1837

Copyright © 2013, Juniper Networks, Inc.

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
   ```

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
   ```

Related Documentation

- Configuring Diameter on page 471

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.

Before you configure Diameter network elements, perform the following task:

- Define the Diameter peers. See "Configuring Diameter Peers" on page 473.

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.

```
[edit diameter network-element dne25 forwarding route dne-route2]
user@host# set function jsrc
```

Related Documentation

- Configuring Diameter on page 471

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 10.9.20.0
```


Copyright © 2013, Juniper Networks, Inc.

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]
   user@host# set routing-instance ri10
   ```

Related Documentation

- Configuring Diameter on page 471

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 *Junos OS System Log Messages Reference*.

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure Diameter base protocol tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the Diameter Base Protocol Trace Log Filename" on page 476.

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

   See "Configuring the Number and Size of Diameter Base Protocol Log Files" on page 476.

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

   See "Configuring Access to the Diameter Base Protocol Log File" on page 477.

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 477.

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

   See "Configuring the Diameter Base Protocol Tracing Flags" on page 478.

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 478.

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 475

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**.

Copyright © 2013, 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. (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**

Related Documentation

- Tracing Diameter Base Protocol Processes for Subscriber Access on page 475

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 475

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 475

Configuring the Diameter Base Protocol Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

  [edit system processes diameter-service traceoptions]
  user@host# **set flag dne**

Related Documentation

- Tracing Diameter Base Protocol Processes for Subscriber Access on page 475

Configuring the Severity Level to Filter Which Diameter Base Protocol Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all** or **verbose**. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages . By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

  [edit system processes diameter-service traceoptions]
  user@host# **set level *severity***

Related Documentation

- Tracing Diameter Base Protocol Processes for Subscriber Access on page 475

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:

 Copyright © 2013, Juniper Networks, Inc.

user@host> **show diameter detail**

- To display summary information about all Diameter instances:

  user@host> **show diameter instance**

- To display detailed information about all Diameter instances:

  user@host> **show diameter instance detail**

- To display information about a specific Diameter instance, add the instance name to the command:

  user@host> **show diameter instance master**

  user@host> **show diameter instance detail master**

- To display summary information about all Diameter routes:

  user@host> **show diameter route**

- To display detailed information about all Diameter routes:

  user@host> **show diameter route detail**

- To display information about a specific Diameter route, add the route name to the command:

  user@host> **show diameter route dne-route2**

  user@host> **show diameter route detail dne-route2**

Related Documentation

- Configuring Diameter on page 471
- Configuring Gx-Plus on page 549
- *Junos OS Operational Mode Commands*

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 ptsp**

- 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**

Related Documentation

- Configuring Diameter on page 471
- Configuring Gx-Plus on page 549
- *Junos OS Operational Mode Commands*

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**

 Copyright © 2013, Juniper Networks, Inc.

- To display detailed statistics about all Diameter peers:

  user@host> **show diameter peer statistics detail**

- To display summary statistics about a specified Diameter peer:

  user@host> **show diameter peer statistics peer235**

- To display detailed statistics about a specified Diameter peer:

  user@host> **show diameter peer statistics detail peer235**

- To delete the specified Diameter peer and all of its statistics.

  user@host>**clear diameter peer peer5 connection**

- To delete the specified Diameter peer and its current statistics:

  user@host>**clear diameter peer peer5 statistics**

Related Documentation

- Configuring Diameter on page 471
- *Junos OS Operational Mode Commands*

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 471
- *Junos OS Operational Mode Commands*

Troubleshooting Diameter Network Configuration

Problem A misconfiguration of the network can prevent Diameter from functioning properly. Configuration options for the Diameter base protocol are simple in the current release, simplifying discovery of a misconfiguration.

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*
- *show diameter network-element map*
- *show diameter peer*
- *show diameter peer map*
- *show diameter route*

Troubleshooting Diameter Network Connectivity

Problem In some circumstances, problems can arise with network connectivity to Diameter peers. The problem may originate with the peer or the peer host.

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.

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *show diameter peer*
- *show diameter peer statistics*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 19

JSRC and Juniper Networks Session Resource Control (SRC) Overview

- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 485
- Service Accounting with JSRC on page 486
- Understanding JSRC-SAE Interactions on page 488

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.

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

Copyright © 2013, Juniper Networks, Inc.

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 488
- Messages Used by Diameter Applications on page 457
- Diameter AVPs and Diameter Applications on page 462
- Configuring JSRC on page 491

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.

 Copyright © 2013, Juniper Networks, Inc.

- Interim accounting—When interim accounting is enabled for a service session, JSRC sends interim accounting messages to the SAE at a specified interval to report the cumulative accounting information available at that time. Interim accounting is ignored when accounting is not enabled for the corresponding service session.

JSRC accounting for a service begins when the service is activated, and remains in effect while the service is active. The SAE specifies the service (policy) to be activated for the subscriber with the Juniper-Policy-Install AVP (AVP code 2020). When this AVP includes the Juniper-Acct-Collect AVP (AVP code 2054), JSRC initiates service activation/deactivation accounting for the service.

JSRC initiates interim accounting when the Juniper-Policy-Install AVP includes the Acct-Interim-Interval AVP (AVP code 85). In this case, JSRC updates the accounting values at the interval specified in the AVP— in the range 600 through 86,400 seconds. Aggregate counters are reported for the dual stack case.

JSRC and the SAE exchange Diameter Accounting-Request (ACR) and Accounting-Answer (ACA) messages to communicate accounting data. Both messages include the Juniper-Acct-Record AVP (AVP code 2053) to identify the service for which accounting information is requested.

JSRC sends ACR messages to report accounting data to the SAE. The ACR message includes the Accounting-Record-Type AVP (AVP code 480) to specify the kind of accounting record that it is sending. When a service is activated, this AVP has a value of START_RECORD. When a service is deactivated, it has a value of STOP_RECORD. For interim accounting, ACR messages are sent at the specified accounting interval and the AVP has a value of INTERIM_RECORD.

In addition to specifying the accounting record type, the ACR messages include standard RADIUS attributes to specify the desired statistics: Acct-Input-Octets [42], Acct-Output-Octets [43], Acct-Input-Packets [47], Acct-Output-Packets [48], and Acct-Session-Time [46].

The SAE returns ACA messages to the JSRC to acknowledge receipt of the ACR messages.

An access profile specifies subscriber access authentication and accounting parameters. When a service is activated through JSRC, the accounting reports can be sent either to the SAE or to RADIUS. The default configuration sends the reports to the SAE; you can also configure this by including the **service accounting-order activation-protocol** statement in the access profile. To send the reports instead to the RADIUS server, include the **service accounting-order radius** statement in the access profile.

When a service is activated through RADIUS rather than through JSRC, the accounting reports of the service session are sent to the RADIUS server.

Related Documentation

- Configuring Service Packet Counting on page 1175
- Messages Used by Diameter Applications on page 457
- Diameter AVPs and Diameter Applications on page 462

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 © 2013, 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 © 2013, 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 485
- Messages Used by Diameter Applications on page 457
- Diameter AVPs and Diameter Applications on page 462
- Configuring JSRC on page 491

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 20

Configuring JSRC for Subscriber Access

- Configuring JSRC on page 491
- Configuring the JSRC Partition on page 492
- Assigning a Partition to JSRC on page 493
- Authorizing Subscribers with JSRC on page 493
- Provisioning Subscribers with JSRC on page 494
- Configuring Service Accounting with JSRC on page 494

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 492.

2. Assign the JSRC partition.

   See "Assigning a Partition to JSRC" on page 493.

3. Configure JSRC authorization for subscribers.

   See "Authorizing Subscribers with JSRC" on page 493.

4. Configure JSRC provisioning for subscribers.

   See "Provisioning Subscribers with JSRC" on page 494.

5. Configure service accounting by JSRC.

   See "Configuring Service Accounting with JSRC" on page 494.

6. Configure JSRC event tracing as part of general authentication service tracing operations.

   See "Tracing General Authentication Service Processes" on page 175.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 485

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 471.

Configuration for the JSRC partition consists of naming the partition and then associating a Diameter instance, the SAE hostname, and the SAE realm with the partition.

To configure the JSRC partition:

1. Create the partition.

   ```
   [edit jsrc]
   user@host# set partition partition1
   ```

2. Specify the Diameter instance for the JSRC partition.

   NOTE: **Currently, only the default Diameter instance, *master*, is supported.**

   ```
   [edit jsrc partition partition1]
   user@host# set diameter-instance master
   ```

3. Configure the destination host for the JSRC partition.

   ```
   [edit jsrc partition partition1]
   user@host# set destination-host sae1
   ```

4. Configure the destination realm for the JSRC partition.

   ```
   [edit jsrc partition partition1]
   user@host# set destination-realm generic.example.com
   ```

Related Documentation

- Configuring JSRC on page 491

Copyright © 2013, Juniper Networks, Inc.

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 492

To assign the JSRC partition:

- Specify the partition name.

```
[edit jsrc]
user@host# set jsrc-partition partition1
```

Related Documentation

- Configuring JSRC on page 491

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 491
- Creating Unique Usernames for DHCP Clients on page 239
- **profile on page 1899**

Copyright © 2013, 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 491

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.

```
[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 486

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 21

Subscribers on Static Interfaces

- Subscribers on Static Interfaces Overview on page 495

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 services static-subscribers login interface *interface-name*** command or the router or process reboots.

Copyright © 2013, Juniper Networks, Inc.

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 © 2013, 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.

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 500
- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 485
- Understanding JSRC-SAE Interactions on page 488

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 22

Configuring Subscribers over Static Interfaces

- Configuring Subscribers over Static Interfaces on page 500
- Specifying the Static Subscriber Global Access Profile on page 501
- Specifying the Static Subscriber Global Dynamic Profile on page 502
- Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers on page 502
- Configuring the Static Subscriber Global Authentication Password on page 503
- Configuring the Static Subscriber Global Username on page 503
- Creating a Static Subscriber Group on page 504
- Specifying the Static Subscriber Group Access Profile on page 505
- Specifying the Static Subscriber Group Dynamic Profile on page 505
- Enabling Multiple Subscribers on a VLAN Logical Interface for a Static Subscriber Group on page 506
- Configuring the Static Subscriber Group Authentication Password on page 506
- Configuring the Static Subscriber Group Username on page 507
- Forcing a Static Subscriber to Be Logged Out on page 508
- Resetting the State of an Interface for Static Subscriber Login on page 508
- Forcing a Group of Static Subscribers to Be Logged Out on page 508
- Resetting the State of an Interface Group for Static Subscriber Login on page 508
- Verifying Information about Subscriber Sessions on Static Interfaces on page 509
- Tracing Static Subscriber Operations on page 509
- Configuring the Static Subscribers Trace Log Filename on page 510
- Configuring the Number and Size of Static Subscribers Log Files on page 510
- Configuring Access to the Static Subscribers Log File on page 511
- Configuring a Regular Expression for Static Subscriber Messages to Be Logged on page 511
- Configuring the Static Subscribers Tracing Flags on page 512

Copyright © 2013, Juniper Networks, Inc.

- Configuring the Severity Level to Filter Which Static Subscriber Messages Are Logged on page 512
- Example: Configuring Static Subscribers for Subscriber Access on page 512

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 501.

2. Specify the global dynamic profile that is instantiated when static subscribers log in.

   See "Specifying the Static Subscriber Global Dynamic Profile" on page 502.

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 502

4. Configure the global authentication password for static subscribers.

   See "Configuring the Static Subscriber Global Authentication Password" on page 503.

5. Configure the global username for static subscribers.

   See "Configuring the Static Subscriber Global Username" on page 503.

6. Configure a group of subscribers to share values different from the global configuration.

   See "Creating a Static Subscriber Group" on page 504.

7. Specify the access profile for the static subscriber group.

   See "Specifying the Static Subscriber Group Access Profile" on page 505.

8. Specify the dynamic profile for the static subscriber group.

   See "Specifying the Static Subscriber Group Dynamic Profile" on page 505.

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 506.

10. Configure the authentication password for the static subscriber group.

 Copyright © 2013, Juniper Networks, Inc.

See "Configuring the Static Subscriber Group Authentication Password" on page 506.

11. Configure the username for the static subscriber group.

    See "Configuring the Static Subscriber Group Username" on page 507.

12. (Optional) Force a static subscriber to be logged out from an interface.

    See "Forcing a Static Subscriber to Be Logged Out" on page 508.

13. (Optional) Enable an interface to accept static subscriber logins.

    See "Resetting the State of an Interface for Static Subscriber Login" on page 508.

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 508.

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 508.

16. Configure trace options for troubleshooting the configuration.

    See "Tracing Static Subscriber Operations" on page 509.

Related Documentation

- Subscribers on Static Interfaces Overview on page 495
- [edit system services static-subscribers] Hierarchy Level on page 1411

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 500
- Specifying the Static Subscriber Group Access Profile on page 505
- **profile on page 1899**

Specifying the Static Subscriber Global Dynamic Profile

You specify a previously created dynamic profile that is instantiated when a static subscriber logs in. This profile is used for all static subscribers. This value can be overridden for a group of static subscribers when a different profile is configured for that group.

To specify the dynamic profile used for all static subscribers:

- Specify the profile name.

```
[edit system services static-subscribers]
user@host# set dynamic-profile dyn-profile-1
```

Related Documentation

- Configuring Subscribers over Static Interfaces on page 500
- Specifying the Static Subscriber Group Dynamic Profile on page 505
- dynamic-profiles on page 1598

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 500
- Specifying the Static Subscriber Group Dynamic Profile on page 505

 Copyright © 2013, Juniper Networks, Inc.

- dynamic-profile on page 1597

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 Gj85*3mS

Related Documentation

- Configuring Subscribers over Static Interfaces on page 500
- Configuring the Static Subscriber Group Authentication Password on page 506
- authentication on page 1468

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.

   [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 500
- Configuring the Static Subscriber Group Username on page 507
- **username-include on page 2101**

Creating a Static Subscriber Group

You can override the configuration that is applied globally to static subscribers by creating a static subscriber group that consists of a set of statically configured interfaces. You can then apply a common configuration for the group with values different from the global values for access and dynamic profiles, password, and username.

To configure an interface group for static subscribers:

1. Access the **[edit system services static-subscribers]** hierarchy level.
2. Create the group and assign the name.

```
[edit system services static-subscribers]
user@host# edit group boston
```

3. Specify the names of one or more interfaces on which static subscribers can be created. You can repeat the **interface *interface-name*** statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

```
[edit system services static-subscribers group boston]
user@host# set interface ge-1/0/1.1
user@host# set interface ge-1/0/1.2
```

4. (Optional) You can use the **upto *upto-interface-name*** option to specify a range of interfaces for a group.

```
[edit system services static-subscribers group boston]
user@host# set interface ge-1/0/1.3 upto ge-1/0/1.9
```

5. (Optional) You can use the **exclude** option to exclude a specific interface or a specified range of interfaces from the group. For example:

```
[edit system services static-subscribers group boston]
user@host# set interface ge-1/0/1.1 upto ge-1/0/1.102
user@host# set interface ge-1/0/1.6 exclude
user@host# set interface ge-1/0/1.70 upto ge-1/0/1.80 exclude
```

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 500
- Specifying the Static Subscriber Group Access Profile on page 505
- Specifying the Static Subscriber Group Dynamic Profile on page 505
- Configuring the Static Subscriber Group Authentication Password on page 506
- Configuring the Static Subscriber Group Username on page 507

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 500
- profile on page 1899

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
```

Related Documentation

- Configuring Subscribers over Static Interfaces on page 500
- Specifying the Static Subscriber Global Dynamic Profile on page 502
- dynamic-profiles on page 1598

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 500
- Specifying the Static Subscriber Group Dynamic Profile on page 505
- **dynamic-profile on page 1597**

Configuring the Static Subscriber Group Authentication Password

You can override the configured global authentication password by specifying a different password for a group of static subscribers. This password is included in the Access-Request message sent to AAA to authenticate all static subscribers in the group.

To specify the authentication password used for a group of static subscribers:

- Specify the password.

  [edit system services static-subscribers group boston authentication]
  user@host# set password knTS$$k2

Related Documentation

- Configuring Subscribers over Static Interfaces on page 500

Copyright © 2013, Juniper Networks, Inc.

- Configuring the Static Subscriber Global Authentication Password on page 503
- authentication on page 1468

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:

2ndfloor.ge-0-1-2-50.default.master.building5.example.com

Related Documentation

- Configuring Subscribers over Static Interfaces on page 500
- Configuring the Static Subscriber Global Username on page 503
- username-include on page 2101

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 508

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 508

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 508

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.

Copyright © 2013, Juniper Networks, Inc.

- 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 508

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 *Junos OS Operational Mode Commands*
- Configuring Subscribers over Static Interfaces on page 500
- Subscribers on Static Interfaces Overview on page 495

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 *Junos OS System Log Messages Reference*.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure static subscriber tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the Static Subscribers Trace Log Filename" on page 510.

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

   See "Configuring the Number and Size of Static Subscribers Log Files" on page 510.

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

   See a"Configuring Access to the Static Subscribers Log File" on page 511.

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 511.

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

   See "Configuring the Static Subscribers Tracing Flags" on page 512.

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 512.

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 509

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

 Copyright © 2013, Juniper Networks, Inc.

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**

Related Documentation

- Tracing Static Subscriber Operations on page 509

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 509

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 509

Configuring the Static Subscribers Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

  [edit system processes static-subscribers traceoptions]
  user@host# **set flag authentication**

Related Documentation

- Tracing Static Subscriber Operations on page 509

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

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all** or **verbose**. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages . By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

  [edit system processes static-subscribers traceoptions]
  user@host# **set level** ***severity***

Related Documentation

- Tracing Static Subscriber Operations on page 509

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 {

 Copyright © 2013, Juniper Networks, Inc.

```
    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 Gj85*3mS;
    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 knTS$$k2;
      username-include {
        user-prefix 2ndFloor;
        interface;
        logical-system-name;
        routing-instance-name;
        domain-name example.net;
      }
    }
  }
}
```

6. Configure tracing options for static subscriber events.

```
static-subscribers {
  traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
```

Related Documentation

- Subscribers on Static Interfaces Overview on page 495
- Configuring Subscribers over Static Interfaces on page 500

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 23

PTSP and Juniper Networks Session and Resource Control (SRC)

- PTSP Overview on page 515
- Juniper Networks Session and Resource Control (SRC) and PTSP Overview on page 516
- Understanding PTSP-SAE Interactions on page 517
- Packet-Triggered Subscribers Services Overview on page 518
- Understanding the Subscriber Profiles for Client Sessions per PSTP Partition on page 522

PTSP Overview

The packet-triggered subscribers and policy control (PTSP) feature allows the application of policies to individual source IP addresses flowing through a given interface. A subscriber context is created for each distinct source IP address seen in a given underlying interface. This feature can be used to support dynamic subscribers that are controlled by a subscriber termination device, such as a B-RAS or GGSN device, that is connected to an MX Series 3D Universal Edge Router.

PTSP has the following responsibilities:

- Create a subscriber context for each distinct IPv4 address on a given interface (subscriber context).
- Apply policies to or remove policies from the subscriber context.
- Collect statistics and report for each individual policy for each subscriber context.
- Derive information about subscribers.

You can associate specific subscriber contexts based on IPv4 addresses and provide service activation and deactivation for these subscribers. The Multiservices DPC (MS-DPC) maintains a table of addresses for each subscriber and any corresponding policies. If an address is not found in the subscriber table, then a new subscriber context is created. All policies are defined on a per-subscriber basis. Once the subscribers are present in the subscriber table, PSTP enforces the policies active for the subscriber context. PTSP can report the subscribers to the SAE using the Diameter protocol so that the SRC software can manage the subscribers and services with dynamic policies. You can also configure static policies, but dynamic policies take precedence over static policies. When you download a new dynamic policy, it takes effect only for new flows. All new flows and

Copyright © 2013, Juniper Networks, Inc.

TCP connections use the new dynamic policy. Existing flows are not affected by the new policy unless they timeout, after which they are considered a new flow.

Statistics collection that is aggregated on a service rule basis is also shared with the SAE using the Diameter application. These statistics are not written to a flat file. Statistics collection that is aggregated on an application or application group basis is written to a flat file. These statistics are not shared with the SAE using the Diameter protocol.

Hardware Requirements for PTSP for Subscriber Access

PTSP is supported on Juniper Networks MX Series 3D Universal Edge Routers. You must have a Multiservices DPC (MS-DPC) on the MX Series router.

Related Documentation

- Configuring PTSP on page 529

Juniper Networks Session and Resource Control (SRC) and PTSP 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 peer on a Juniper Networks routing platform and the remote SRC peer on a C Series Controller. The local peer is known as PTSP 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.

The SRC software enables the SAE to activate and deactivate subscriber services (described by SRC policies). The SAE installs or removes policies using a service rule policy template called **__svc_rule__**. This policy template indicates which policy is applied to a new subscriber session. Additional policies are bound to new sessions; they do not affect existing sessions. Note that policy name must be unique between PPR requests. You can use the same rule name within a single request, but you cannot use the same name again in a separate request.

Statistics collection that is aggregated on a service rule basis is also shared with the SAE using the Diameter protocol.

NOTE: **More than one Diameter-based application (function) can run on a router simultaneously.**

Copyright © 2013, Juniper Networks, Inc.

NOTE: When the SRC software downloads PTSP policies, it matches all the application groups defined in the rule of the PTSP policy if the **appliation-group-any** keyword is used in the policy. The **application-group-any** keyword is not configured on the router although the application group name is defined in the application identification configuration database on the router to process application-aware access list (AACL) rules for accepting or discarding packets. The keyword is considered as an exception because the application group is defined in the application identification database.

Related Documentation

- Messages Used by Diameter Applications on page 457
- Diameter AVPs and Diameter Applications on page 462
- Understanding PTSP-SAE Interactions on page 517
- Configuring the PTSP Application on page 525
- Configuring PTSP on page 529

Understanding PTSP-SAE Interactions

This topic describes the sequences of Diameter messages exchanged between PTSP and the SAE as they interact to perform the following tasks for subscriber access:

- Subscriber login

  When a packet-triggered subscriber logs in, PTSP sends a Diameter AA-Request message to request service provisioning from the SAE that includes the Session-Id attribute for the new subscriber. If the AA-Request fails, then the subscriber is not considered logged in and the subscriber session is not managed by the SAE. Only the static PTSP rules apply to the subscriber.

  The SAE returns a Diameter AA-Answer message with the Result-Code. The AA-Answer message can include the Juniper-Policy-Install AVP (AVP code 2020), which is used to specify a service to attach to the subscriber's IP address.

  PTSP can send an AA-Request message to the SAE to confirm activation. The SAE returns a AA-Answer message in acknowledgment. If the AA-Request message fails or the SAE does not respond with an AA-Answer message, the subscriber session is managed by the SAE.

- Service activation and deactivation

  The SAE policies provision subscriber services. After a packet-triggered subscriber is logged in, the SAE can send a PPR message to PTSP to activate or deactivate services. A given PPR can include the Juniper-Policy-Install AVP (AVP code 2020) to activate a service or the Juniper-Policy-Remove AVP (AVP code 2027) to deactivate a service.

  PTSP 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.

Copyright © 2013, Juniper Networks, Inc.

- Resynchronization

  Either PTSP 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.

  PTSP initiates resynchronization at startup, such as when AAA starts or restarts. PTSP uses the Juniper-Last-Origin-Host AVP (AVP code 2055) to keep track of the active SAE host in a multi-SAE environment. When an SAE in a multi-SAE environment becomes active, it must send an SRQ to PTSP as its first message. PTSP initiates a synchronization when it receives any other message type from an SAE that is different from the SAE indicated in the Juniper-Last-Origin-Host AVP.

  Both entities initiate a resynchronization by sending an SRQ message. The recipient responds with an SRR message.

- 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 and Accounting-Answer messages include the Juniper-Acct-Record AVP (AVP code 2053) which identifies the policy for which accounting information is requested.

- Subscriber logout

  PTSP can determine when there is a logout request for a packet-triggered subscriber in two ways:

  - The SAE terminates a subscriber session by sending an ASR message to PTSP.
  - PTSP monitors a subscriber session and starts the logout process after 30 minutes of inactivity.

  The subscriber logout triggers the final statistics aggregation for all policies and the removal of any policies installed by the SAE. PTSP sends an STR message that indicates the logout event to the SAE.

Related Documentation

- Juniper Networks Session and Resource Control (SRC) and PTSP Overview on page 516
- Messages Used by Diameter Applications on page 457
- Diameter AVPs and Diameter Applications on page 462
- Configuring the PTSP Application on page 525
- Configuring PTSP on page 529

Packet-Triggered Subscribers Services Overview

The packet-triggered subscribers and policy control (PTSP) feature allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device. You can associate specific subscriber contexts based on IPv4 addresses and provide dynamic service activation and deactivation for these subscribers. Once the subscribers

Copyright © 2013, Juniper Networks, Inc.

are present in the subscriber database on the router, PSTP can report the subscribers to the SAE using the PTSP application so that the SRC software can manage the subscribers and services.

PTSP policies can be downloaded dynamically from the external policy manager (such as SRC) or configured statically on the router. The PTSP policies can be configured for each distinct IPv4 source address for a given interface on which the service is configured. Each distinct IPv4 address is considered a subscriber and all PTSP policies are applied on a per-subscriber basis. Dynamic policies, which are always specific to a subscriber, take precedence over static policies.

You can set up PTSP policies to:

- Manage traffic by configuring filtering, rate-limiting, and QoS enforcement in the rules.
- Steer traffic by specifying the forwarding instance in the forward rule.
- Collect accounting information by service rule or by application.

When you configure PTSP policies, you must specify the type of statistics collection (**count**) and the IP address used to identify the packet-triggered subscriber (**demux**) in the service rule. All service rules attached to a given service set must have the same settings for these options.

For the statistics collection type, terms and rules also cannot mix and match the following styles:

- rule—Statistics are aggregated in one bucket for the service rule and Diameter is used to report the statistics.
- application—Statistics are aggregated by application for a specific application, for a specific application group, or in one bucket. The statistics are reported in a flat file.

Subscriber instantiation is triggered for ingress packets by the IP address. When source address is specified, the source IP address of the ingress packets is used to establish the subscriber context. When destination address is specified, the destination IP address of the ingress packets is used to establish the subscriber context. If the IP address does not correspond to a known subscriber, then a new subscriber context is created to log in the packet-triggered subscriber.

The match conditions include local address, local port, remote address, and remote port. The following table describes how the **demux** value changes the IP address or port used for these terms.

Match Conditions	demux source-address		demux destination-address	
	Ingress Flows	Egress Flows	Ingress Flows	Egress Flows
local-address	Source address	Destination address	Destination address	Source address
remote-address	Destination address	Source address	Source address	Destination address
local-port	Source port	Destination port	Destination port	Source port

Match Conditions	demux source-address		demux destination-address	
	Ingress Flows	Egress Flows	Ingress Flows	Egress Flows
remote-port	Destination port	Source port	Source port	Destination port

Subscriber Identification Method for PTSP Partition

The PSTP functionality uses RADIUS attributes, such as *User-Name* to identify subscribers in a RADIUS partition. If a service provider uses a different RADIUS attribute other than User-Name, the authentication of subscribers and establishment of client sessions fail. To enable service providers to use a subscriber-identification method that suits their network needs, you can add flexible configurations in the packet-triggered subscriber process.

The PTSP configurable user-identification feature allows you to do the following:

- Configure the subscriber identification method for PTSP partitions, based on the network topology and the service provider requirements.
- Insert subscriber-specific tags for the subscriber's HTTP traffic for which the reference to subscriber-specific tagging is provided using subscriber identification.

The PTSP application generates the subscriber-identification parameter as a text-string by combining the RADIUS attribute value and the internal attribute value of the PTSP partition. The text-string is generated in the same order as the attributes that are configured in the PTSP partition.

NOTE: **Only RADIUS partitions support user-identification to configure the subscriber-identification method for PTSP partitions.**

PTSP Services on Aggregated and Redundant Services PICs

The packet-triggered subscribers and policy control (PTSP) feature supports both Aggregated Multiservices (AMS) and Redundant Multiservices (RMS) PICs. RMS services interfaces support 1:1 redundancy between two logical PICs and in an active or standby model. AMS services interfaces support load sharing and N:1 redundancy between N logical PICs.

NOTE: **The PTSP services do not support load balancing on AMS.**

In 1:1 redundancy, if services PIC fails:

- The subscriber is logged out, the traffic is switched to the redundant services PIC, and the subscriber receives a new session ID to log in with.
- The subscriber's last configured accounting data is retrieved as the latest interim accounting record.

Copyright © 2013, Juniper Networks, Inc.

In AMS, the PTSP subscriber's traffic is redistributed to other services PIC and the same subscriber may appear on different services PICs. The subscriber with no new data flow is logged out after idle timeout with the complete accounting data. The following example depicts the AMS scenario:

```
ams0 {
  load-balancing-options {
    member-interface mams-4/0/0;
    member-interface mams-4/1/0;
    member-interface mams-5/0/0;
    member-failure-options {
      redistribute-all-traffic;
    }
  }
  unit 1 {
    family inet;
  }
}
```

The traffic on ms-4/0/0 is redistributed to ms-4/1/0 only after ms-5/0/0 has failed. In this example, there are two subscribers: s1 on ms-4/0/0 and s2 on ms-4/1/0. The two subscribers have the same source IP address. If there is no new traffic, s1 is eventually logged out after idle timeout.

NOTE: PTSP does not support any type of hash key for traffic sharing among logical PICs configured with the same PTSP service set. For PTSP to work, all traffic for any given subscriber needs to reach the same logical PIC within an AMS container. For this to happen, the AMS hashing algorithm needs to align with the PTSP demux type, as follows:

- If PTSP is configured for source-demux, then the AMS hashing algorithm must be based on the source-ip-address only.
- If PTSP is configured for destination-demux, then the AMS hashing algorithm must be based on the destination-ip-address only.
- No other type of AMS hashing algorithm is compatible with PTSP.

NOTE: The packet level idle timeout for every packet is assigned from a given subscriber transiting the router. If the timeout limit sets in, the subscriber is logged out. The valid range for the subscriber packet idle timeout is 15 to 1440 minutes.

Related Documentation

- Configuring PTSP on page 529
- Configuring Static PTSP Rules on page 531

Understanding the Subscriber Profiles for Client Sessions per PSTP Partition

Subscriber profiles for service activation enables you to specify which service plug-ins become activated on a per-subscriber basis. Previously, the only control mechanism for specifying service activation was to attach a service-set configuration to a selected interface or route. The new utility allows you to enable or disable services based on the subscriber associated with every data flow. As a result, you can apply differentiated services to different sets of subscribers. You can exercise the control mechanism in one of two ways: by using a CLI operational command or a RADIUS attribute.

NOTE: **This feature applies only to MP-SDK services and does not depend on the specific services enabled or disabled, except that PTSP must be included in the chain.**

The procedure consists of three steps:

1. Configure a service set that includes all the services to be applied to flows. You can include a default subscriber profile that controls which services are and are not active by default. The default profile applies to all subscribers until overridden for a specific subscriber. In the absence of a default subscriber profile, all services specified in the service set are applied by default. You can also include one or more alternative subscriber profiles that can be implemented to override the default profile. The following sample configuration illustrates these components:

```
services {
  service-set ss1 {
    application-identification-profile appidr1;
    idp-profile idpr1;
    aacl-rules aaclr1;
    hcm_rules hcmr1;
    sfw_rules sfwr1;
    subscriber-profile {
      sp1;
    }
    interface-service {
      service-interface ms-3/0/0.0;
    }
  }
  }
  subscriber-profile sp1 {
    disable HCM;
    enable IDP {
      concurrent-data-sessions 10;
    }
    disable AACL;
    max-data-sessions-per-subscriber {
      limit 10;
      exceed-action [ syslog drop ];
    }
  }
  subscriber-profile sp2 {
```


Copyright © 2013, Juniper Networks, Inc.

```
        enable HCM;
        disable IDP;
        enable AACL,
        max-data-sessions-per-subscriber {
          limit 100;
          exceed-action [ syslog ];
        }
      }
```

Initially, all traffic reaching the service plane under service set ss1 receives all the services configured in service set ss1 that are enabled by the default subscriber profile sp1 applied to it. In the example, APPID, stateful firewall, and IDP are enabled, whereas HCM and AACL are disabled. However IDP is enabled for only at most 10 sessions concurrently. Beyond that threshold, IDP is also disabled. Also, because of the max-data-sessions-per-subscriber setting, any subscriber is allowed a maximum of ten concurrent data sessions. Beyond that threshold, data sessions are logged and dropped.

2. There are two ways to dynamically override the default subscriber profile associated with a particular PTSP subscriber:

   - CLI operational command
   - RADIUS attribute or VSA in an access-accept message.

   From the previous example, assume that the subscriber profile for subscriber X is dynamically set to sp2. After that, any new data session associated with subscriber X has a different set of services applied to it. In the example, it would be APPID, stateful firewall, HCM, and AACL. Also, because the max-data-sessions-per-subscriber setting changes to 100, subscriber X now has no upper limit on the number of concurrent data sessions, although if that number crosses the 100 threshold, the threshold-crossing event is logged.

   The following examples illustrate the dynamic override settings:

   Operational command

   ```
   user@router>request services subscriber clear subscriber-profile
   client-id client-id
   ```

   ```
   user@router>request services subscriber set subscriber-profile
   subscriber-profile-name client-id client-id
   ```

   RADIUS configuration

   ```
   user@router# set system services packet-triggered-subscribers
   partition-radius foo subscriber-service-profile attribute-26.4874.31
   ```

3. Processing of a new data session at the service plane takes place as follows, with respect to subscriber profiles:

   1. A new flow starts. MP-SDK sends a SESSION-INTEREST event to the service plug-ins. The first plug-in in the chain is the subscribers (PTSP) plug-in.
   2. The subscribers plug-in matches the flow to its subscriber by searching its database. It sets the subscriber ID in the session metadata.

3. The subscriber plug-in checks for the corresponding subscriber profile and which services are enabled. It then sets the services mask of enabled and disabled services in the session metadata.
4. MP-SDK or JSF invokes only the services that are enabled per the services mask. The other services are skipped, even if configured in the service set.

NOTE: : Subscriber-profile changes affect only the upcoming flows. Existing flows remain unaffected.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 24

Configuring the PTSP Application

- Configuring the PTSP Application on page 525
- Configuring the PTSP Partition on page 526
- Assigning the PTSP Partition on page 526
- Tracing Packet-Triggered Subscriber Operations on page 527

Configuring the PTSP Application

You can configure the PTSP client application to work with the Session and Resource Control (SRC) peer to centrally manage packet-triggered subscribers and services. PTSP 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. The PTSP application also performs statistics collection and reporting.

To configure the PTSP application:

1. Configure the PTSP partition.

   See "Configuring the PTSP Partition" on page 526.

2. Assign the PTSP partition.

   See "Assigning the PTSP Partition" on page 526.

3. Configure statistics collection and reporting.

   See "Tracing Packet-Triggered Subscriber Operations" on page 527.

Related Documentation

- Juniper Networks Session and Resource Control (SRC) and PTSP Overview on page 516

Configuring the PTSP Partition

PTSP works within a specific logical system:routing instance context, called a partition. The partition is configured to connect to the external policy manager.

NOTE: **Currently, only a single partition is supported; you must configure it within the default logical system:routing instance context.**

Before you configure the PTSP partition to connect to the external policy manager, perform the following task:

- Configure the Diameter instance for the remote SRC peer at the **[edit diameter]** hierarchy level. See "Configuring Diameter" on page 471.

Configuration for the PTSP 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 PTSP partition:

1. Create the partition at the **[edit system services packet-triggered-subscribers]** hierarchy level.

   [edit system services packet-triggered-subscribers]
   user@host# **edit partition ptsp–default**

2. Specify the Diameter instance for the PTSP partition.

   [edit system services packet-triggered-subscribers partition ptsp–default]
   user@host# **set diameter-instance master**

3. Configure the destination host for the PTSP partition.

   [edit system services packet-triggered-subscribers partition ptsp–default]
   user@host# **set destination-host sae1**

4. Configure the destination realm for the PTSP partition.

   [edit system services packet-triggered-subscribers partition ptsp–default]
   user@host# **set destination-realm generic.example.com**

5. Configure the subscriber ID for the PTSP partition.

   [edit system services packet-triggered-subscribers partition-radius
   radius-partition-name]
   user@host# **set subscriber-identification**

Related Documentation

- Configuring the PTSP Application on page 525

Assigning the PTSP Partition

You must associate the PTSP partition with the logical system:routing instance.

Copyright © 2013, Juniper Networks, Inc.

NOTE: Currently, only the global logical system:routing instance, *master* logical system and default routing instance, is supported.

Before you assign the PTSP partition, perform the following task:

- Configure the PTSP partition. See "Configuring the PTSP Partition" on page 526.

To assign the PTSP partition:

- Specify the partition name at the **[edit system]** hierarchy level.

```
[edit system]
user@host# set packet-triggered-subscribers-partition ptsp–default
```

Related Documentation

- Configuring the PTSP Application on page 525

Tracing Packet-Triggered Subscriber Operations

Packet-triggered subscriber tracing operations track packet-triggered subscriber operations and record them in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

All log files are located in the **/var/log** directory. You cannot change the directory (**/var/log**) in which trace files are located. When the trace file reaches its maximum size, a **.0** is appended to the filename, then a new file is created with a **.1**, and finally a **.2**. When the maximum number of trace files is reached, the oldest trace file is overwritten.

To configure packet-triggered subscriber tracing operations:

1. Specify that you want to configure tracing options.

```
[edit system services packet-triggered-subscribers]
user@host# edit traceoptions
```

2. (Optional) Configure the name for the file used for the trace output.
3. (Optional) Configure the number and size of the log files.
4. (Optional) Configure flags to filter the operations to be logged.

The packet-triggered subscriber traceoptions operations are described in the following sections:

- Configuring the Packet-Triggered Subscribers Trace Log Filename on page 528
- Configuring the Size of Packet-Triggered Subscribers Log Files on page 528
- Configuring the Packet-Triggered Subscribers Tracing Flags on page 528

Copyright © 2013, Juniper Networks, Inc.

Configuring the Packet-Triggered Subscribers Trace Log Filename

By default, the name of the file that records trace output for packet-triggered subscribers is **jptspd**. You can specify a different name with the **file** option.

To configure the filename for packet-triggered subscribers tracing operations:

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

```
[edit system services packet-triggered-subscribers traceoptions]
user@host# set file ptsp-subs_1
```

Configuring the Size of Packet-Triggered 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 size of trace files:

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

```
[edit system services packet-triggered-subscribers traceoptions]
user@host# set file ptsp-subs_1 _logfile_1 size 2097152
```

Configuring the Packet-Triggered Subscribers Tracing Flags

To configure the flags for the events to be logged:

- Configure the flags.

```
[edit system services packet-triggered-subscribers traceoptions]
user@host# set flag peer
user@host# set flag session
```

Related Documentation

- Configuring the PTSP Application on page 525

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 25

Configuring Packet-Triggered Subscriber Services

- Configuring PTSP on page 529
- Configuring the Multiservices DPC for PTSP on page 530
- Configuring PTSP Service Rules on page 531
- Configuring Static PTSP Rules on page 531
- Configuring PTSP Rule Sets on page 533
- Configuring PTSP Service Sets on page 534
- Configuring the PTSP Forwarding Instance on page 534
- Configuring a Statistics Profile for PTSP on page 536
- Tracing PTSP Operations on page 538
- Verifying and Managing PTSP Configuration on page 539

Configuring PTSP

You can configure the packet-triggered subscribers and policy control (PTSP) feature on MX Series routers to allow the application of policies to dynamic subscribers that are controlled by a subscriber termination device, such as a B-RAS or GGSN device, connected to an MX Series router. The subscribers are associated by their IPv4 address and dynamic or static policies can be applied. Dynamic policies take precedence over static policies. When you download a new dynamic policy, it takes effect only for new flows. All new flows and TCP connections use the new dynamic policy. Existing flows are not affected by the new policy unless they timeout, after which they are considered a new flow.

To configure PTSP services on the MX Series router:

1. Configure the Multiservices DPC.

   See "Configuring the Multiservices DPC for PTSP" on page 530.

2. Configure the Diameter application to support the download of dynamic PTSP policies from the external policy manager (such as SRC). The PTSP application also provides statistics collection and reporting.

   See "Configuring the PTSP Application" on page 525.

3. Configure the static PTSP service rules.

   See "Configuring Static PTSP Rules" on page 531.

4. Configure statistics collection and reporting in a flat file.

   See "Configuring a Statistics Profile for PTSP" on page 536 and "Tracing PTSP Operations" on page 538.

Related Documentation

- PTSP Overview on page 515

Configuring the Multiservices DPC for PTSP

To configure the Multiservices Dense Port Concentrator (MS-DPC) to support PTSP services, perform the following tasks:

- Enabling the PTSP Service Package on the Multiservices DPC on page 530
- Configuring Services Interface for PTSP on page 530

Enabling the PTSP Service Package on the Multiservices DPC

The PTSP feature runs on the Multiservices DPC, you must enable the PTSP service package on the Multiservices DPC before you can configure the PTSP software. The name of the PTSP service package is **jservices-ptsp**.

To enable the PTSP service package:

1. Determine the FPC slot number and the PIC number of the MS-DPC on which you want to enable the PTSP service package.

   ```
   user@host> show chassis hardware
   ```

   In this example, the FPC slot number is 3 and the PIC number is 0.

2. Enable the jservices-ptsp package on the Multiservices DPC.

   ```
   [edit chassis]
   user@host# set fpc 3 pic 0 adaptive-services service-package extension-provider
     package jservices-ptsp
   ```

Configuring Services Interface for PTSP

NOTE: **ams- interfaces and rms- interfaces can be configured for PTSP.**

To configure the services interface for PTSP:

1. Enter edit mode for the interface.

   ```
   [edit]
   user@host# edit interfaces ms-3/0/0
   ```

2. Configure a logical unit and specify the protocol family.

   ```
   [edit interfaces ms-3/0/0]
   ```

 Copyright © 2013, Juniper Networks, Inc.

```
user@host# set unit 0 family inet
```

Related Documentation

- Configuring PTSP on page 529
- PTSP Overview on page 515

Configuring PTSP Service Rules

PTSP policies can be downloaded dynamically from the external policy manager (such as SRC) or configured statically on the router. The PTSP policies can be configured for each distinct IPv4 source address for a given interface on which the service is configured. Each distinct IPv4 address is considered a subscriber and all PTSP policies are applied on a per-subscriber basis.

Dynamic policies, which are always specific to a subscriber, take precedence over static policies. When you download a new dynamic policy, it takes effect only for new flows. All new flows and TCP connections use the new dynamic policy. Existing flows are not affected by the new policy unless they timeout, after which they are considered a new flow.

To configure the PTSP policies, perform these tasks:

- To download dynamic policies and to collect statistics with Diameter, configure the Diameter application for PTSP. See "Configuring the PTSP Application" on page 525.
- To configure static policies, see "Configuring Static PTSP Rules" on page 531. To collect statistics in a flat file, see "Configuring a Statistics Profile for PTSP" on page 536.

Related Documentation

- Configuring PTSP on page 529
- PTSP Overview on page 515

Configuring Static PTSP Rules

You can configure the static PTSP policies on the router. If the PTSP service is configured on the underlying interface, the PTSP service enforces the policies associated with the subscriber context.

To configure static PTSP rules:

1. Specify the rule that you want to configure.

   ```
   [edit services ptsp]
   user@host# edit rule ptspRule1
   ```

2. Specify the direction in which the rule match is applied.

   ```
   [edit services ptsp rule ptspRule1]
   user@host# set match-direction input
   ```

3. Specify the IP address used for the subscriber context. Subscriber instantiation is always triggered for ingress packets, so this value indicates which IP address in the ingress packets for the flow is used.

   [edit services ptsp rule ptspRule1]
   user@host# **set demux source-address**

4. Specify the statistics aggregation, collection, and reporting style. Terms and rules cannot mix and match different styles.

   [edit services ptsp rule ptspRule1]
   user@host# **set count-type rule**

   If you specify the rule style, statistics collection is performed by the Diameter application. If you specify the application style, statistics collection is in a flat file controlled by the local policy decision function (L-PDF).

5. (Optional) Specify the forward rule used for forwarding packets. See "Configuring the PTSP Forwarding Instance" on page 534.

   [edit services ptsp rule ptspRule1]
   user@host# **set forward-rule *forward-rule-name***

6. Configure the term precedence for the rule.

   [edit services ptsp rule ptspRule1]
   user@host# **edit term 1**

7. Configure the match conditions for the term. See Table 51 on page 532.

   [edit services ptsp rule ptspRule1 term 1]
   user@host# **set from remote-address-range low 203.0.0.2 high 203.0.0.100**
   user@host# **set from remote-address-range low 204.0.0.2 high 204.0.0.253**

8. (Optional) Specify the action taken when the match conditions are met. See Table 52 on page 533.

   [edit services ptsp rule ptspRule1 term 1]
   user@host# **set then count rule**
   user@host# **set then accept**

Table 51 on page 532 describes the match conditions for PTSP rules.

Table 51: PTSP Match Conditions

Match Condition	Description
application-group-any	Application group name defined in the application identification configuration.
application-groups [***application-group-name***]	Application group name defined in the application identification configuration.
applications	Application name defined in the application identification configuration.
local-port-range low ***low-value*** **high** ***high-value***	Local port range.
local-ports ***value-list***	Local ports.

 Copyright © 2013, Juniper Networks, Inc.

Table 51: PTSP Match Conditions *(continued)*

Match Condition	Description
protocol ***protocol-number***	IP protocol number.
remote-address (*address* | any-unicast)	Remote IP address. IPv4 only.
remote-address-range low ***low-value*** **high** ***low-value***	Remote address range. IPv4 only.
remote-port-range low ***low-value*** **high** ***high-value***	Remote port range.
remote-ports ***value-list***	Remote ports.
remote-prefix-list ***prefix-list-name***	Prefixes in the specified list.

Table 52 on page 533 describes the actions for PTSP rules.

Table 52: PTSP Actions

Action or Action Modifier	Description
accept	Accept the packet.
count	Increment the specified counter.
discard	Drop the packet.
forwarding-class	Classify the packet into the specified forwarding class.
police	Rate-limit packets based on the specified policer.

Related Documentation

- Configuring the PTSP Forwarding Instance on page 534
- Configuring a Statistics Profile for PTSP on page 536
- Configuring PTSP on page 529
- PTSP Overview on page 515
- Packet-Triggered Subscribers Services Overview on page 518

Configuring PTSP Rule Sets

You can define a collection of PTSP rules to determine the actions performed on packets.

To configure static PTSP rule sets:

1. Specify the rule set that you want to configure.

   [edit services ptsp]

```
user@host# edit rule-set ptspRules
```

2. Specify the rules in the order that you want them processed.

```
[edit services ptsp rule-set ptspRules]
user@host# set rule ptspRule1
user@host# set rule ptspRule2
```

Related Documentation

- Configuring Static PTSP Rules on page 531

Configuring PTSP Service Sets

To configure the service set for the PTSP application:

1. Configure the service set that you want to contain the PTSP service.

```
[edit services service-set ptspServiceSet]
user@host# set service-set ptspServiceSet
```

2. Specify the PTSP rules that constitute the service set that is applied to the services interface.

```
[edit services service-set ptspServiceSet]
user@host# set ptsp-rules ptsp-rule1
user@host# set ptsp-rules ptsp-rule2
```

3. Configure the services interface.

NOTE: ams- interfaces and rms- interfaces are supported for PTSP.

```
[edit services service-set ptspServiceSet]
user@host# set interface-service service-interface ms-3/0/0.0
```

4. Associate the service set with the underlying interface from which the subscribers originate. The service set must be applied to the interface facing the subscriber, that is, the interface with the IP address of the subscriber.

```
[edit interfaces ge-4/0/0 unit 0 family inet service]
user@host# set input service-set ptspServiceSet
user@host# set output service-set ptspServiceSet
```

Related Documentation

- Configuring Static PTSP Rules on page 531
- Configuring PTSP Rule Sets on page 533

Configuring the PTSP Forwarding Instance

Before you can forward PTSP traffic, perform these tasks for each forwarding instance:

1. Configure each PTSP forwarding instance as a routing instance type of forwarding.
2. Configure a firewall filter with an action that specifies the routing instance configured in Step 1.

 Copyright © 2013, Juniper Networks, Inc.

3. Configure the unit number for the Multiservices interface that specifies the filter configured in Step 2 as the input filter.

   NOTE: To avoid service set dependency on specific unit numbers, use the same unit number across all Multiservices interfaces where PTSP services are applied.

4. Configure the PTSP forward rule to specify the forwarding instance.

   NOTE: When the forwarding instance action is performed on the flow, any postservice filters are not applied to the underlying interface.

If you want to forward traffic for PTSP subscribers, you must specify the forwarding instance for specific subscribers based on IP address, network, or prefix list. The match direction for forward rules is always input.

To configure the PTSP forwarding instance:

1. Specify the PTSP forward rule that you want to use when configuring a PTSP forwarding instance.

   ```
   [edit services ptsp]
   user@host# edit forward-rule ptspForward
   ```

2. Set the term precedence for the forward rule. Term with lowest precedence is evaluated first.

   ```
   [edit services ptsp forward-rule ptspForward]
   user@host# edit term 5
   ```

3. Configure the match conditions for the IP address, address range, or prefix list. See Table 53 on page 535.

   ```
   [edit services ptsp forward-rule ptspForward term 5]
   user@host# set from local-address 200.0.0.1
   ```

Table 53: PTSP Forward Rule Match Conditions

Match Condition	Description
application-groups [*application-group-name*]	Application group name defined in the application identification configuration.
applications	Application name defined in the application identification configuration.
local-address (*address* | **any-unicast**)	Local IP address. IPv4 only.
local-address-range low *low-value* **high** *high-value*	Local address range. IPv4 only.
local-prefix-list *prefix-list-name*	Prefixes in the specified list.

NOTE: You can specify match conditions for applications or application groups that support application identification (APPID) services, but we do not recommend specifying the forwarding instance action when you are using these match conditions in PTSP policies. In this situation, some network topologies may route packets in a manner that causes the flow to be dropped. For example, the APPID services might forward some packets on the default routing instance while the PTSP services forward other packets in the same flow to another routing instance.

4. Configure the forwarding instance action with the routing instance name and the unit number.

```
[edit services ptsp forward-rule ptspForward term 5]
user@host# set then forwarding-instance less-effort-ri 144
```

NOTE: When the forwarding instance action is performed on the flow, any postservice filters are not applied to the underlying interface.

Related Documentation

- For information about APPID services, see the *Junos OS Services Interfaces Library for Routing Devices*
- For information about forwarding instances, see the *Junos OS Routing Protocols Library for Routing Devices*

Configuring a Statistics Profile for PTSP

The local policy decision function (L-PDF) enables you to configure properties for statistics output by creating a statistics profile. The statistics profile configures the files to which statistics records are exported and the format that is exported. You configure the statistics profile so that the statistics records are exported to a flat file. Flat files contain statistics that are collected for each subscriber by application or application group. The statistics in a flat file are not transmitted to the external policy manager using Diameter.

To configure a statistics profile for PTSP:

1. Specify that you want to configure a statistics profile.

```
[edit system services local-policy-decision-function]
user@host# edit statistics
```

2. Configure the file properties used for the trace output.
3. Configure the profile properties.
4. Specify the record type.

Copyright © 2013, Juniper Networks, Inc.

Tasks to configure a statistics profile for PTSP are:

- Configuring the File Properties for Statistics Data Output on page 537
- Configuring the Profile Properties for Statistics Data Output on page 537
- Configuring the Record Type for Statistics Data on page 538

Configuring the File Properties for Statistics Data Output

You configure a file to which the statistics data output is exported in a specified format.

To configure the file properties:

1. Specify the unique filename for receiving statistics data output.

   [edit system services local-policy-decision-function statistics]
   user@host# **edit file ptsp**

2. (Optional) Specify the maximum number of files that are maintained at one time and the maximum size of each file. If you configure one of these options, you also must set the other option.

   [edit system services local-policy-decision-function statistics file ptsp]
   user@host# **set files 10 size 1g**

3. Specify the interval for transferring files to archive sites.

   [edit system services local-policy-decision-function statistics file ptsp]
   user@host# **set transfer-interval 60**

4. Specify one or more URLs for archiving the files. Archiving can be done by using FTP or SCP.

   [edit system services local-policy-decision-function statistics file ptsp]
   user@host# **set archive-sites "ftp://anonymous@10.227.1.114"**

Configuring the Profile Properties for Statistics Data Output

You can create an AACL statistics profile, which configures the statistics to collect and write to a file in the **/var/stats/aacl** directory.

To configure the profile properties:

1. Specify the name of the profile.

   [edit system services local-policy-decision-function statistics]
   user@host# **edit aacl-statistics-profile ptsp**

2. (Optional) Specify the file in the **/var/stats/aacl** directory in which statistics are collected. Enclose the name within quotation marks.

   [edit system services local-policy-decision-function statistics aacl-statistics-profile ptsp]
   user@host# **set file "pstp"**

3. Set the interval for reporting statistics.

   [edit system services local-policy-decision-function statistics aacl-statistics-profile ptsp]

```
user@host# set report-interval 5
```

4. Set the **interim-active-only** mode for reporting statistics. This mode reports only statistics that have changed in the past report interval.

```
[edit system services local-policy-decision-function statistics aacl-statistics-profile
  ptsp]
user@host# set record-mode interim-active-only
```

5. Specify the statistics to be collected in the log file.

```
[edit system services local-policy-decision-function statistics aacl-statistics-profile
  ptsp]
user@host# set aacl-fields all-fields
```

Configuring the Record Type for Statistics Data

You must configure the interim record type for recording the AACL statistics.

To configure the record type:

- Specify interim as the record type.

```
[edit system services local-policy-decision-function statistics]
user@host# set record-type interim
```

Related Documentation

- Tracing PTSP Operations on page 538
- Configuring PTSP on page 529

Tracing PTSP Operations

Tracing operations track L-PDF operations and record them in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

By default, no events are 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, **ptspd**. 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 *Junos OS System Log Messages Reference*.)

Copyright © 2013, Juniper Networks, Inc.

To customize trace file settings:

1. Specify that you want to configure tracing options.

   [edit system services local-policy-decision-function]
   user@host# **edit traceoptions**

2. Configure the filename used for the trace output.

   [edit system services local-policy-decision-function traceoptions]
   user@host# **set file lpdfd**

3. (Optional) Configure the maximum number and size of the log files. If you configure one of these options, you also must set the other option.

   [edit system services local-policy-decision-function traceoptions]
   user@host# **set files 10 size 1g**

4. (Optional) Specify flags to filter the operations to be logged. To specify more than one flag, include multiple **flag** statements.

   [edit system services local-policy-decision-function traceoptions]
   user@host# **set flag ptsp-statistics**

   The following table describes the flags that you can include.

Flag	Description
configuration	Trace configuration events
database	Trace database events
general	Trace general flow
ptsp-statistics	Trace PTSP events
rtsock	Trace routing socket events
statistics	Trace statistics events
subscriber	Trace subscriber events

Related Documentation

- Configuring a Statistics Profile for PTSP on page 536
- Configuring PTSP on page 529

Verifying and Managing PTSP Configuration

Purpose Display and clear information about packet-triggered subscribers and PTSP services.

Action

- To display bandwidth information about subscribers:

  user@host> **show services subscriber bandwidth**

- To display information about the active dynamic policies applied to a subscriber:

user@host> **show services subscriber dynamic-policies client-id *client-id***

- To display information about the data flows associated with a subscriber:

  user@host> **show services subscriber flows client-id *client-id***

- To display information about the active packet-triggered subscriber sessions on the router:

  user@host> **show services subscriber sessions**

- To display information about the data traffic statistics for the packet-triggered subscriber:

  user@host> **show services subscriber statistics client-id *client-id***

- To clear the active packet-triggered subscriber session on the router and log out the subscriber:

  user@host> **clear services subscriber sessions client-id *client-id***

Related Documentation

- *Junos OS Operational Mode Commands*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 26

Gx-Plus Overview

- Gx-Plus for Provisioning Subscribers Overview on page 541
- Understanding Gx-Plus Interactions Between the Router and the PCRF on page 543

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 54 on page 541 are interchangeable.

Table 54: Differences Between Gx-Plus and Junos OS Terminology

Gx-Plus	Junos OS
policy	service
rule	service
rule install or installation	service activation or instantiation
rule uninstall	service deactivation
usage monitoring	service accounting

Gx-Plus enables the router acting as a PCEF to exchange Diameter Credit-Control Application (DCCA) messages with a PCRF residing on a server to request credit authorization and service provisioning for authenticated subscribers. When an application requests AAA to activate a subscriber's session, the router sends a Credit-Control-Request

Copyright © 2013, Juniper Networks, Inc.

(CCR) message to determine whether the subscriber has credit for the desired services and to request provisioning of those services from the PCRF policy manager.

The PCRF responds with a Credit-Control-Answer (CCA) message that indicates success or failure for credit authorization. If the subscriber has sufficient credit for the requested services, credit is authorized. If the subscriber has insufficient credit for the services, credit authorization fails.

The CCA can include services to be activated for the subscriber. If the response times out, the subscriber is logged in but only default services—if present—are activated for the subscriber. The router interprets the omission of the Result-Code AVP from the CCA as a provisioning authorization failure and does not allow the subscriber to log in.

When a subscriber client application, such as DHCP, sends a subscriber logout notice to AAA, the router in turn sends a CCR message to the PCRF to request subscriber termination. The PCRF acknowledges the logout with a CCA message.

Different Diameter message types exchanged by the router and the PCRF contain different sets of attribute-value pairs (AVPs). If data for an AVP is not available for a request to the PCRF, that AVP is omitted from the message. If the PCRF subsequently has insufficient information to decide on the request, it may deny the request.

Gx-Plus establishes sessions that correspond to IPv4 DHCP sessions on dual-stack IPv6/IPv4 or IPv4-only subscriber interfaces, depending on the access profile. By default, IPv6 information is not communicated to the PCRF. You must explicitly configure Gx-Plus to include IPv6 information. When you do so, Gx-Plus can establish sessions that correspond to DHCPv6 sessions on IPv6-enabled subscriber interfaces and on dual-stack IPv6/IPv4-enabled interfaces.

For dual-stack DHCP subscribers (DHCPv4 and DHCPv6 on the same VLAN), each DHCP session is treated as a separate Gx-Plus session. However, only a single Gx-Plus session exists for dual-stack PPP sessions.

Gx-Plus includes the following fault tolerance and recovery capabilities:

- Unlimited retries of unacknowledged provisioning requests
- Unlimited retries of logout requests
- Router event notification
- Router discovery

NOTE: **More than one Diameter-based application (function), such as Gx-Plus or JSRC, can run on a router simultaneously.**

Related Documentation

- Messages Used by Diameter Applications on page 457
- Diameter AVPs and Diameter Applications on page 462
- Understanding Gx-Plus Interactions Between the Router and the PCRF on page 543

Copyright © 2013, Juniper Networks, Inc.

• Configuring Gx-Plus on page 549

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) 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

Copyright © 2013, Juniper Networks, Inc.

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.
   - 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.

 Copyright © 2013, Juniper Networks, Inc.

a. If the service statistics meet or exceed any of these trigger thresholds during the session, the router sends a CCR-U message to the PCRF with accounting information in the Usage-Monitoring-Information AVP (AVP code 1067). The AVP now contains the Used-Service-Unit AVP (AVP code 446) to report the current values for all four statistics.

b. In response, the PCRF may return a CCA-U message with the Usage-Monitoring-Information AVP, which can include any of the following: the Granted-Service-Unit AVP with new threshold triggers (absolute values rather than increments to the previous thresholds), the Charging-Rule-Install AVP (AVP code 1001) for service activations, or the Charging-Rule-Remove AVP (AVP code 1002) for service deactivations.

NOTE: **The router does not aggregate statistics across services.**

6. When the subscriber logs out, the router sends a CCR-T message (termination notice) to the PCRF, which responds with a CCA-T message.

Fault Tolerance and Event Notification

Although the probability is low, the PCRF and the router can have different values for the number of subscribers. This error can arise from the following scenarios:

- CCA-I loss: if no CCA-I is delivered to the router, then the PCRF considers a subscriber as provisioned whereas the router considers it not provisioned.
- CCR-T loss: if no CCR-T is delivered to the PCRF, then the PCRF considers a subscriber to be provisioned whereas the router considers the subscriber not provisioned (logged out).

Loss of messages can be greater during cold boots and high availability events. Unacknowledged CCR-I and CCR-T requests are retransmitted forever until a satisfactory response is received to reduce the incidence of failure, and significant events are reported to Gx-Plus. By default, the number of outstanding requests is limited to 40 to avoid overloading the PCRF. This limit reduces the possibility of losing requests. You can modify this number by including the **max-outstanding-requests** statement at the **[edit access-gx-plus global]** hierarchy level.

Gx-Plus does not rely on the connection state between devices to detect router or PCRF outages, because some events do not affect the connection state and others are not detected when there is a Diameter relay or proxy between the devices. Event notifications (JSER messages) are sent when certain events take place on the router. The Juniper-Event-Type AVP (AVP code 2103) in the message describes the event.

Event notifications are retried until Gx-Plus returns a JSEA message with a Result-Code value of DIAMETER_SUCCESS (2001) to acknowledge receipt of the event notification. When retrying notifications, one notification is sent for each outstanding event. No other request are sent as long as there is any outstanding event other than an application watch dog (AWD).

Copyright © 2013, Juniper Networks, Inc.

Table 55 on page 546 lists router events and the subsequent router and PCRF actions.

Table 55: Router Events, Router Actions, and PCRF Actions

Router Event	Router Action	PCRF Action
The router receives no response from the PCRF or an error response.	Send event notification.	Respond to event notification.
The configuration changes. Significant changes such as the origin host or realm and the Gx-Plus partition destination host or realm also increment the value of the Origin-State-Id AVP.	Send event notification.	Respond to event notification and perform discovery.
The router receives an explicit discovery request from the PCRF.	Send event notification.	Respond to event notification.
The router undergoes a cold boot and all sessions are lost. This can result from a catastrophic failure or power cycle.	Send event notification.	Respond to event notification and clear the database.
The router undergoes a warm boot.	Send event notification.	Respond to event notification and clear the database.
Recovery resources that are needed to continuously retry unacknowledged requests (CCR-N and CCR-T messages) are exhausted. The value of the Origin-State-Id AVP is incremented. This event is unlikely to occur.	Send event notification.	Respond to event notification and perform discovery.

An important aspect of Gx-Plus fault tolerance is that subscriber login and termination requests are retried (replayed) forever until a satisfactory response is received from the PCRF. In rare circumstances, this can result in a stack of pending requests being replayed over and over.

You can issue the **clear network-access gx-plus replay** command to clear all pending requests. This command causes Gx-Plus to send a JSER message to PCRF that includes the Juniper-Event-Type AVP (AVP code 2103) with a value of 3 indicating a discovery request. The PCRF then returns a JDER message to initiate discovery of all subscribers. When this discovery completes, all pending subscriber requests are cleared.

PCRF-Generated Discovery

The PCRF runs a discovery process in response to data loss, exhaustion of router resources, operator request, or router request. The JSDR message specifies the level of verbosity desired in the reply from Gx-Plus. The message also specifies whether the request is for data about a particular session or information similar to an SNMP Get-Bulk for all sessions. Gx-Plus returns a JSDA message that indicates complete success, limited success, or an error. In the event of success, the requested data is also returned.

 Copyright © 2013, Juniper Networks, Inc.

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.

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.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Gx-Plus for Provisioning Subscribers Overview on page 541
- Messages Used by Diameter Applications on page 457
- Diameter AVPs and Diameter Applications on page 462
- Configuring Gx-Plus on page 549
- Default Subscriber Service Overview on page 1127
- Configuring a Default Subscriber Service on page 1127

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 27

Configuring Gx-Plus

- Configuring Gx-Plus on page 549
- Configuring the Gx-Plus Partition on page 550
- Configuring Gx-Plus Global Attributes on page 551
- Provisioning Subscribers with Gx-Plus on page 552
- Disabling PCRF Control of a Subscriber Session on page 552

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 550.

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 551.

3. Configure Gx-Plus provisioning for subscribers.

   See "Provisioning Subscribers with Gx-Plus" on page 552.

4. (Optional) Override PCRF control of a subscriber session to correct services or troubleshoot a problem.

   See "Disabling PCRF Control of a Subscriber Session" on page 552.

5. (Optional) Configure Gx-Plus event tracing as part of general authentication service tracing operations.

Copyright © 2013, Juniper Networks, Inc.

See "Tracing General Authentication Service Processes" on page 175.

Related Documentation

- Gx-Plus for Provisioning Subscribers Overview on page 541

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 471.

Configuration for the Gx-Plus partition consists of naming the partition and then associating a Diameter instance, the PCRF hostname, and the PCRF realm with the partition.

To configure the Gx-Plus partition:

1. Create the partition or specify the name of an existing partition.

   [edit access gx-plus]
   user@host# set partition *partition-name*

2. Specify the Diameter instance for the Gx-Plus partition.

   NOTE: **Currently, only the default Diameter instance, *master*, is supported.**

   [edit access gx-plus partition *partition-name*]
   user@host# set diameter-instance *instance-name*

3. (Optional) Configure the destination host for the Gx-Plus partition.

   [edit access gx-plus partition *partition-name*]
   user@host# set destination-host *hostname*

4. Configure the destination realm for the Gx-Plus partition.

   [edit access gx-plus partition *partition-name*]
   user@host# set destination-realm *realm*

The following example shows a Gx-Plus partition configuration.

```
gx-plus {
  partition partition1 {
    diameter-instance master;
    destination-host pcrf1;
    destination-realm generic.example.com;
  }
```

 Copyright © 2013, Juniper Networks, Inc.

```
}
```

Related Documentation

- Configuring Gx-Plus on page 549
- Gx-Plus for Provisioning Subscribers Overview on page 541

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
```

Related Documentation

- Configuring Gx-Plus on page 549
- Gx-Plus for Provisioning Subscribers Overview on page 541

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 149.

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 549
- Gx-Plus for Provisioning Subscribers Overview on page 541

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**

 Copyright © 2013, Juniper Networks, Inc.

2. (Optional) Verify that provisioning is disabled for the session.

   user@host> **show network-access aaa subscribers session-id *subscriber-session-id* detail**

For example, to disable provisioning for subscriber larry:

```
user@host> show network-access aaa subscribers
Username       Logical system/Routing instance   Client type    Session-ID
...
larry          default:default                   dhcp           55
...

user@host> request network-access aaa subscriber set session-id 55 provisioning-state none
user@host> show network-access aaa subscribers session-id 55 detail
Type: dhcp
Username: larry@isp5.net
Stripped username: larry
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 1122

Copyright © 2013, Juniper Networks, Inc.

PART 8

Mobile IP Access

- Mobile IP Overview on page 557
- Configuring Mobile IP on page 569

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 28

Mobile IP Overview

- Mobile IP Home Agent Elements and Behavior on page 557
- Mobile IP Registration on page 560
- Mobile IP Routing and Forwarding on page 564
- Mobile IP in the WiMAX Environment on page 565

Mobile IP Home Agent Elements and Behavior

Mobile IP is a tunneling-based solution that enhances the utility of Junos routing platforms at the edge of the network between fixed wire and wireless network domains. This tunneling-based solution enables a router on a user's home subnet to intercept and forward IP packets to users who roam beyond traditional network boundaries. Mobile IP is useful in environments where mobility is desired and the traditional land line dial-in model does not provide an adequate solution, and in environments where a wireless technology is used.

You configure Mobile IP home agent parameters in the **[edit services mobile-ip]** hierarchy level, the **[edit logical-systems *logical-system-name*]** hierarchy level, and the **[edit routing-instances *routing-instances-name*]** hierarchy level.

NOTE: **Currently, Junos OS does not support configuration of the Mobile IP foreign agent.**

Traditionally, IP addresses are associated with a fixed network location. To achieve mobility, the mobile node assumes a secondary IP address that matches the new network and redirects the traffic bound to the primary or home address to the mobile node's new network. In the Mobile IP architecture, the two agents that accomplish this task are the home agent and the foreign agent.

When a mobile node roams into a new, foreign network, it negotiates with the foreign agent to get a secondary IP address, which is referred to as the care-of address. The mobile node registers this care-of address with the home agent. The home agent then establishes a tunnel to the care-of address if the tunnel is not established earlier.

Copyright © 2013, Juniper Networks, Inc.

NOTE: You need to establish only one tunnel between the home agent and the care-of address. Demultiplexing of the traffic is done through IP address inspection.

Packets sent to the home address of the mobile node are redirected by the home agent through the tunnel to the care-of address at the foreign agent. The foreign agent routes the packets to the mobile node's home address. Figure 10 on page 558 illustrates this forwarding and routing process behavior. Although the traffic to the correspondent node comes from the foreign agent, to the correspondent node the traffic appears to come from the mobile node's home network.

Figure 10: Mobile IP Network Without Reverse Tunneling

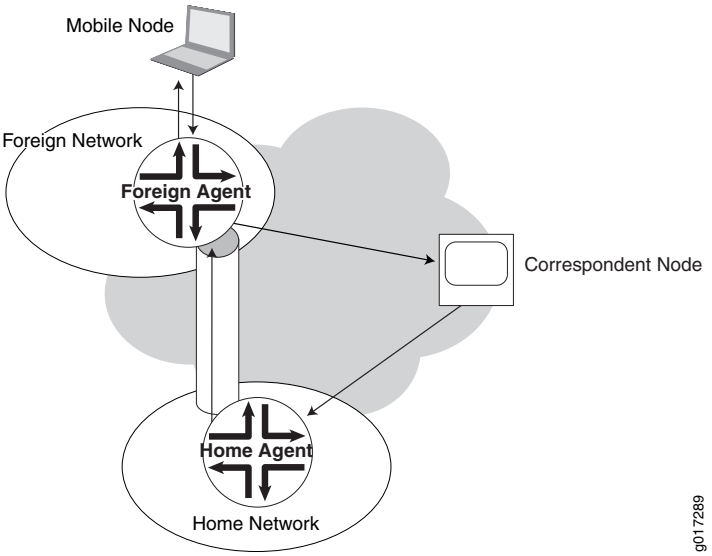

If the mobile node's home address is a private address or if the foreign agent implements ingress filtering, a reverse tunnel from the care-of address to the home agent is required. This reverse tunnel capability is negotiated between the foreign agent and the home agent when the mobile node requests registration. Traffic from a correspondent node to the mobile node is forwarded by the home agent through the foreign agent as in the other scenario. Figure 11 on page 559 shows how traffic from the mobile node to a correspondent node is tunneled from the foreign agent to the home agent and then routed to the correspondent node by the home agent.

Copyright © 2013, Juniper Networks, Inc.

Figure 11: Mobile IP Network with Reverse Tunneling

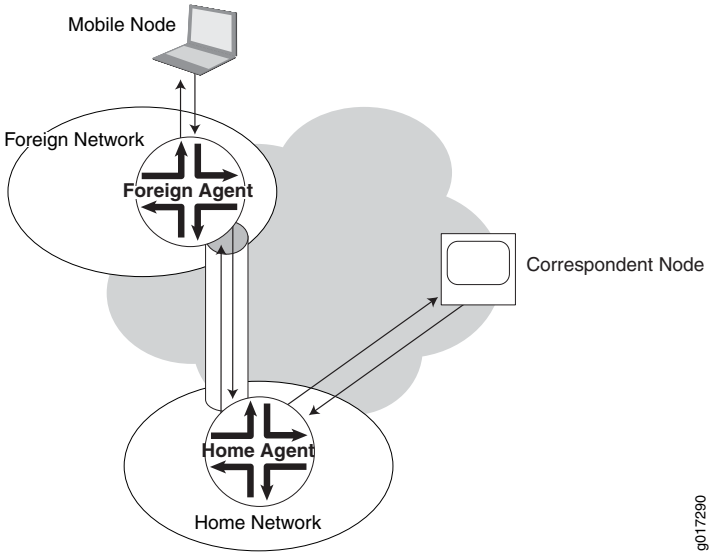

Mobile nodes typically belong to a virtual network, which is an address range or subnet that is not directly served by any physical, routed interface on the home network. These mobile nodes never return home to attach to a physical interface on the home agent. Traffic destined for the mobile node can be forwarded over any interface.

You can use the Mobile IP home agent feature to configure the home agent within the default router context with either local or AAA authentication. When you configure local authentication, you can also configure Mobile IP independently in any named routing instance in any configured logical router. When you configure AAA as the authentication method, you can configure Mobile IP only in the default router context.

The Mobile IP home agent can also receive, process, and send Worldwide Interoperability for Microwave Access (WiMAX) vendor-specific RADIUS attributes (VSAs). This feature enables Mobile IP home agent to work in a WiMAX home connectivity services network (H-CSN), to provide for mobility management at the IP layer.

The home agent handles the following tasks:

- Registration of mobile nodes
- Routing and forwarding of mobile node traffic

Related Documentation

- Mobile IP Registration on page 560
- Mobile IP Routing and Forwarding on page 564
- Mobile IP in the WiMAX Environment on page 565
- Configuring Mobile IP on page 569

Mobile IP Registration

The home agent receives the registration requests (RRQs) on UDP port 434. The registration request contains the home agent IP address. The home agent can support static home address allocation and dynamic home address allocation. The home agent can revoke a mobile node's registration. When this happens, the mobility binding is removed and the foreign agent is informed of the revocation so it can free up its resources. The foreign agent can send a registration revocation request to the home agent when the mobile node roams to another area. The revocation request can include a revocation support extension to indicate that it supports the revocation mechanism.

Home Address Assignment

The mobile node's home address can either be preconfigured, or dynamically allocated by the Mobile IP home agent. If a nonzero home address is preconfigured, the home agent processes the registration request using the home address and NAI (if the NAI is present).

If the home address is dynamically allocated, the mobile node submits a zero home address and requests the home agent to assign an IP address. The mobile node then uses the address provided by the home agent for subsequent registration requests, until the mobile node is rebooted or the registration expires.

Home address allocation is done by one of the existing authentication, authorization, and accounting (AAA) server back-end address mechanisms, such as:

- By RADIUS, in the Framed-IP-Address attribute
- From a local address pool returned by RADIUS in the Framed-Pool attribute

Authentication

The home agent authenticates the requests based on RFC 3344—IP Mobility Support for IPv4 (August 2002). By default, a AAA server is used for authentication; alternatively, you can configure local authentication parameters on the home agent. The mobile node authentication is verified and the authentication algorithm and key are retrieved by checking the security association indexed by the security parameter index (SPI) value. This verification results in the key and the authentication algorithm with which to compute an MD-5 message digest over the registration request. The Mobile IP home agent supports both HMAC-MD5 and keyed-MD5 authentication algorithms. When the result of this computation matches the authenticator, the mobile-home extension is authenticated. For local authentication, the key is limited to a maximum of 128 bits. For AAA authentication, the key can be longer depending on the maximum length configured on the AAA server.

When HA receives the access accept from the AAA, it extracts the MN-HA key from the response. The home agent does the MN-HA authentication extension processing based on the MN-HA key by running authentication algorithm (HMAC-MD5 or Keyed-MD5) on the message to compute a hash (authenticator), which is compared with the hash value in the MN-HA extension. If the hash value matches, the RRQ is considered authenticated.

 Copyright © 2013, Juniper Networks, Inc.

If a security association is configured for the foreign agent, the foreign-home authentication extension is verified; otherwise, authentication success is based only on the mobile-home authenticator.

The home agent checks the identification (ID) field to verify that a registration message has been freshly generated by the mobile node, and is not simply being replayed by an attacker from some previous registration. The ID field represents a 64-bit Network Time Protocol (NTP)-formatted time value. The configured replay timestamp defines the tolerance time window in seconds by which a registration request timestamp and the local time of the HA can differ. By default, the timestamp must be within 7 seconds of the replay tolerance configured for the mobile node or, if that is configured, the timestamp tolerance of the home agent itself.

Reauthentication

Reauthentication is not currently supported by the authentication process. Mobile IP caches a security association for each mobile node, which helps overcome this limitation. When a mobile node requests re-registration or de-registration, Mobile IP refers to the cached security association for that mobile node and performs MD5 message authentication.

When the security association for the mobile node changes after the node is authenticated, the cache entry is not invalidated. Consequently, the mobile node's RRQ is rejected. In this case you must clear the binding with the mobile node so that it can de-register and then log in.

RADIUS server configuration changes relating to the subscriber do not propagate to the cache. In this case you must clear the binding with the mobile node so that it can de-register and then log in.

AAA Authentication

You can store the security associations and configuration information remotely on a RADIUS server. The home agent applies the authentication algorithm and security key to the mobile node's message. The AAA server uses Juniper Networks vendor-specific attributes (VSAs; vendor ID 4874) listed in Table 56 on page 561. These VSAs are mandatory in the reply to provide the appropriate authentication algorithm and the secure key for the authentication request. If the security parameters are not retrieved, then the request for mobility service is rejected, a security violation error is logged, and no registration reply is generated.

Table 56: Juniper Networks VSAs Used by Mobile IP

Attribute Number	Attribute Name	Description	Value
26–84	Mobile-IP-Algorithm	Authentication algorithm used for Mobile-IP registration	integer: 4-octet
26–85	Mobile-IP-SPI	Security parameter index for Mobile IP registration	integer: 4-octet

Copyright © 2013, Juniper Networks, Inc.

Table 56: Juniper Networks VSAs Used by Mobile IP *(continued)*

Attribute Number	Attribute Name	Description	Value
26–86	Mobile-IP-Key	Security association MD5 key for Mobile IP registration	string: key
26–87	Mobile-IP-Replay	Replay timestamp for Mobile IP registration	integer: 4-octet
26–89	Mobile-IP-Lifetime	Registration lifetime for Mobile IP registration	integer: 4-octet

AAA authentication is accomplished by generating a AAA access-request to a AAA server. This is the default authentication mode, but you can include the **authenticate order aaa** statement at the **[edit services mobile-ip]** hierarchy level to explicitly configure AAA authentication. You cannot configure a fallback mechanism for AAA authentication. If the AAA request times out, the home agent does not fall back on the local router to determine the authentication parameters. The registration request is rejected. When the message is authenticated, the AAA server always returns either the Framed-IP-Address or Framed-Pool attribute for the user.

The presence of the mobile node's NAI and home IP address in the authentication request that the home agent sends to the AAA server is determined by their presence in the mobile node RRQ received by the home agent:

- When both the NAI and home IP address of the mobile node are present in the registration request, then the authentication request from Mobile IP to AAA has the NAI as the user name.
- When only the NAI is present in the registration request, then the NAI is used as the user name.
- When only the IP address (home address) is present in the registration request, then the IP address is used as the user name.
- When both the NAI address and the IP address are missing from the registration request, then the registration request is rejected.

Local Authentication

As an alternative to the default authentication by AAA server, you can store the security associations and configuration information locally on the router hosting the home agent. Local authentication is accomplished by querying the locally configured security parameters for the mobile node. The home agent applies the authentication algorithm and security key to the mobile node's message. If the security parameters are not available or do not match the RRQ, then the request for mobility service is rejected, a security violation error is logged, and no registration reply is generated.

For local authentication, include the **authenticate order local** statement at the **[edit services mobile-ip]** hierarchy level. You cannot configure a fallback mechanism for local

 Copyright © 2013, Juniper Networks, Inc.

authentication. If the local authentication fails, the home agent does not fall back on the AAA server to determine the authentication parameters. The registration request is rejected. Include the **peer** statement at the **[edit services mobile-ip]** hierarchy level to configure the authentication attributes on the home agent for a user identified by IP address or network address identifier (NAI). This user can be a mobile node or a foreign agent.

The authentication attributes include a security parameter index (SPI) to identify a particular security context between the home agent and the mobile node or foreign agent among the contexts available in the mobility security association. Associated with each SPI is the MD5 algorithm and key used to authenticate messages from the mobile node or foreign agent. You can also configure the replay timestamp tolerance for the mobile node or foreign agent.

When local authentication is configured, you can configure Mobile IP independently in any named routing instance in any configured logical router. All Mobile IP statements are available in those contexts, except for the **order aaa** statement at the **[edit services mobile-ip authenticate]** hierarchy level.

Accounting

The Junos Mobile IP home agent application supports time-based accounting for Mobile IP subscribers. Include the **statistics time** statement in the subscriber access profile at the **[edit access profile profile-name accounting]** hierarchy level. Time-based accounting for Mobile IP subscribers also requires that you include the **authenticate order aaa** statement at the **[edit services mobile-ip]** hierarchy level. Accounting begins when the Mobile IP home agent registers the mobile node and creates a binding with the mobile node.

Accounting stops when the binding is deleted. Any of the following actions can cause the binding to be deleted:

- The mobile user logs off.
- The binding lifetime expires.
- The mobile node is deregistered for any reason.
- The foreign agent sends a revocation message.

The Acct-Start message the home agent sends to the AAA server includes the network address identifier (NAI) in the User-Name attribute and the home address of the mobile IP node in the Framed-IP-Address attribute. The Acct-Stop message additionally includes the Acct-Session-Id and Acct-Session-Time attributes.

You cannot currently configure time-based accounting for only the Mobile IP service in a given logical router or routing instance. Enabling time-based accounting for Mobile IP also enables time-based accounting for all other services that are configured in that logical router or routing instance. If you do not want time-based accounting to apply to other services, then you must configure those services in a different logical router or routing instance.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- For information about the specific Juniper Networks VSAs used for Mobile IP RADIUS-based authentication, see Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- Mobile IP Home Agent Elements and Behavior on page 557
- Mobile IP Routing and Forwarding on page 564
- Mobile IP in the WiMAX Environment on page 565
- Configuring Mobile IP on page 569

Mobile IP Routing and Forwarding

Mobile IP employs a care-of address to process traffic for the mobile node.

The mobile node acquires the a care-of address from the foreign agent. The care-of address is reachable from the mobile node, and routable from the home agent. The mobile node includes the care-of address in its registration request to the home agent. After AAA or local authentication successfully processes and authenticates the RRQ and provides both the authorization parameters for the mobile node and an IP address, the home agent then sets up the data path for the mobile node and sends back a registration reply (RRP) confirming successful registration of the mobile node.

When the foreign agent receives the successful RRP from the home agent, the foreign agent sets up the data path for the mobile node. Then it sends the RRP to the mobile node to acknowledge that the mobile node is now successfully registered and the data path between the home agent and the mobile node is in place.

The home agent supports generic routing encapsulation (GRE) and IP-in-IP tunnel encapsulation for forward and reverse tunneling. The tunnels must be statically configured. When packets destined for the mobile node reach a home agent, the home agent encapsulates the packets and tunnels them to the care-of address. Packets that exceed the maximum transmission unit (MTU) value of the tunnel are dropped and an ICMP error message is sent to the source IP address. Packets without an access route are returned to the source with an ICMP destination unreachable error message. For reverse tunnels, packets are de-tunneled and forwarded towards the next hop to the destination address.

Mobile IP does not support graceful Routing Engine switchover (GRES). It handles the rebooting of processes in the following ways:

- Mobile IP process—After Mobile IP completes a restart, it removes the Mobile IP subscriber entries from AAA and the session database. When that is complete, Mobile IP can process new mobile node registration requests.
- AAA process—After AAA completes a restart, Mobile IP removes all subscriber data held internally by AAA and all corresponding session database entries.
- Routing protocol process—When the connection between the routing protocol process and Mobile IP is lost, Mobile IP responds by clearing the mobile node bindings that are associated with the logical system in which the routing protocol process restarted. The

 Copyright © 2013, Juniper Networks, Inc.

routing protocol process maintains routes to mobile nodes during the restart. The routing protocol process flushes these routes if they are not reinstalled after the restart completes and before the stale route timer expires.

Related Documentation

- Mobile IP Home Agent Elements and Behavior on page 557
- Mobile IP Registration on page 560
- Mobile IP in the WiMAX Environment on page 565
- Configuring Mobile IP on page 569

Mobile IP in the WiMAX Environment

Worldwide Interoperability for Microwave Access (WiMAX) is the international standard for wide area radio access networks. It provides a framework for networks that are implemented in different ways to successfully interoperate with mobile subscribers that roam among the networks. This interoperability enables the subscribers to be authenticated by their home network wherever they roam, and to receive the services for which they are authorized.

The Mobile IP home agent can operate in either of two access modes, generic and WiMAX. The generic access type is appropriate when the home agent is deployed in a generic Mobile IP home network. When deployed as a home agent in a WiMAX home connectivity services network (H-CSN), you must configure the WiMAX access type. The WiMAX access type enables the Mobile IP home agent to receive, process, and send WiMAX vendor-specific attributes (VSAs) that are used by AAA and the RADIUS server to authenticate the mobile subscriber. When the access type is generic, the Mobile IP home agent cannot handle these VSAs.

NOTE: **The Mobile IP configuration for WiMAX requires that AAA be used for the authentication method. For that reason, WiMAX is available only in the default router context.**

A WiMAX H-CSN is analogous to the Mobile IP home network for non-WiMAX implementations. When WiMAX is enabled for the Mobile IP home agent in an H-CSN, the Mobile IP home agent triggers subscriber authentication when the agent receives the registration request. The home agent stores WiMAX Forum (vendor ID 24757) vendor-specific attributes (VSAs) listed in Table 57 on page 566 in the session database based on the registration request.

Table 57: WiMAX Forum VSAs used by Mobile IP

Attribute Number	Attribute Name	Description	Value
26–1	WiMAX-Capability	Identifies the WiMAX capabilities supported by the home agent (sent in the Access-Request message). In an Access-Accept message, Identifies the capabilities selected by the RADIUS server (returned in the Access-Accept message).	string or integer
26–6	hHA-IP-MIP4	IP address of the home home agent (hHA) making the request	octet string: IP address
26–10	MN-HA-MIP4-KEY	MN-hHA key sent by the RADIUS server for validation by the home agent	integer: 2-octet salt followed by 16-octet encrypted MN-hHA hash key
26–11	MN-HA-MIP4-SPI	Security parameter index (SPI) associated with the MN-HA-MIP4 key	integer: 4-octet
26–15	hHA-RK-KEY	Key used by the NAS to generate FA-HA keys	integer: 2-octet salt followed by 16-octet encrypted MN-hHA hash key
26–16	hHA-RK-SPI	SPI associated with the hHA-RK key	integer: 4-octet
26–17	HA-RK-Lifetime	Lifetime of the hHA-RK key and derived keys	integer: 4-octet
26–18	RRQ-HA-IP	IP address of the home agent contained in the Mobile IP registration request or the binding update	octet string: IP address

Copyright © 2013, Juniper Networks, Inc.

Table 57: WiMAX Forum VSAs used by Mobile IP *(continued)*

Attribute Number	Attribute Name	Description	Value
26–19	RRQ-MN-HA-KEY	The MN-HA key bound to the home agent IP address as reported by the RRQ-HA-IP attribute. Used to validate the MN-HA-AE of the Mobile IP registration request.	integer: 2-octet salt followed by 16-octet encrypted MN-hHA hash key

The home agent requests AAA to fetch the corresponding WiMAX-related information from the RADIUS server. The AAA client sends an Access-Request message to the server. The RADIUS server responds with the necessary WiMAX information, such as the MN-HA key and the HA-RK key, and then the AAA client passes the response to the home agent. The Mobile IP home agent verifies the response received from AAA, processes the registration request, and then grants, extends, or denies subscriber registration.

Figure 12 on page 567 shows the elements of a sample WiMAX topology.

Figure 12: Sample Mobile IP WiMAX Topology

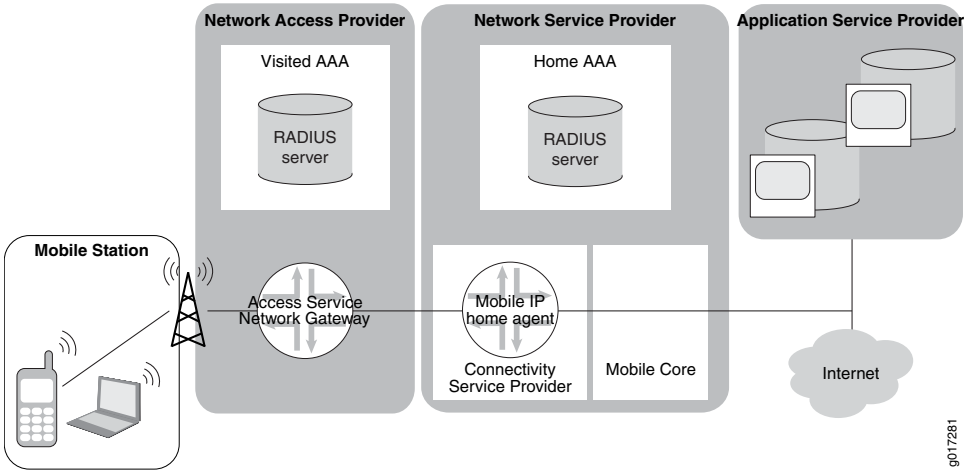

The Mobile IP subscriber registration flow is a four-step process.

1. The access service network gateway (ASN-GW) sends the subscriber registration request from the mobile node to the Mobile IP home agent. The registration request is protected by the MN-HA authentication extension and the FA-HA authentication extension.
2. The home agent requests that the RADIUS server send the cryptographic keys for the Mobile IP session identified by user@realm. The home agent notifies the RADIUS server that it seeks to source IP session-based accounting messages.

Copyright © 2013, Juniper Networks, Inc.

3. The RADIUS server agrees to use IP session-based accounting, provides the requested cryptographic keys, and sends the AAA-Session-ID for this session.
4. The home agent replies to the Mobile IP registration request.

Reauthentication of WiMAX subscribers is not currently supported.

You can configure the Mobile IP home agent for WiMAX access by including the **wimax** statement at the **[edit services mobile-ip access-type**] hierarchy level. You can prevent the Mobile IP home agent from being able to process WiMAX VSAs by either removing the **wimax** statement at the **[edit services mobile-ip access-type**] hierarchy level or by including the **generic** statement at the **[edit services mobile-ip access-type**] hierarchy level. The default access type for Mobile IP home agent is generic.

Related Documentation

- For information about the specific Juniper Networks VSAs used for Mobile IP RADIUS-based authentication, see Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- Mobile IP Home Agent Elements and Behavior on page 557
- Mobile IP Registration on page 560
- Mobile IP Routing and Forwarding on page 564
- Configuring Mobile IP on page 569

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 29

Configuring Mobile IP

- Configuring Mobile IP on page 569
- Configuring the Mobile IP Authentication Method on page 570
- Configuring the Mobile IP Home Agent on page 570
- Configuring the Local Authentication Attributes for the Mobile Node on page 571
- Configuring Accounting for Mobile IP Subscribers on page 572
- Configuring Dynamic Home Assignment for the Mobile Node on page 572
- Configuring the Access Type for Mobile IP on page 573
- Tracing Mobile IP Operations for Subscriber Access on page 573
- Configuring the Mobile IP Trace Log Filename on page 575
- Configuring the Number and Size of Mobile IP Log Files on page 575
- Configuring Access to the Mobile IP Log File on page 576
- Configuring a Regular Expression for Mobile IP Messages to Be Logged on page 576
- Configuring the Mobile IP Tracing Flags on page 576
- Configuring the Severity Level to Filter Which Mobile IP Messages Are Logged on page 577

Configuring Mobile IP

You can configure Mobile IP to provide mobility for subscribers in IP networks. The Mobile IP home agent authenticates registration requests from mobile users and forward traffic to them at their care-of address without having to advertise that address to the wider network.

To configure Mobile IP for mobile subscriber access:

1. Configure the authentication method for registration requests, local or AAA.

   See "Configuring the Mobile IP Authentication Method" on page 570.

2. Configure the Mobile IP home agent.

   See "Configuring the Mobile IP Home Agent" on page 570.

3. Configure the authentication attributes for the mobile node.

   See "Configuring the Local Authentication Attributes for the Mobile Node" on page 571.

Copyright © 2013, Juniper Networks, Inc.

4. Configure accounting for Mobile IP subscribers.

   See "Configuring Accounting for Mobile IP Subscribers" on page 572

5. Configure the dynamic reassignment of the mobile node to another home agent.

   See "Configuring Dynamic Home Assignment for the Mobile Node" on page 572.

6. Configure the access type for Mobile IP.

   See "Configuring the Access Type for Mobile IP" on page 573.

7. Configure trace options for troubleshooting the configuration.

   See "Tracing Mobile IP Operations for Subscriber Access" on page 573.

Configuring the Mobile IP Authentication Method

You can configure Mobile IP to authenticate registration requests from mobile nodes by either the locally configured attributes or a AAA server. AAA server authentication is the default method.

NOTE: AAA server authentication is available only in the default router context. Local authentication is available in both default and nondefault router contexts.

To configure the Mobile IP authentication method:

- Specify either local or AAA authentication.

  ```
  [edit services mobile-ip]
  user@host# set authenticate order local
  ```

Related Documentation

- Configuring Mobile IP on page 569

Configuring the Mobile IP Home Agent

To configure the home agent for a Mobile IP virtual network:

1. Configure the loopback IP address that is used as the home agent IP address.

   ```
   [edit services mobile-ip home-agent virtual-network]
   user@host# set home-agent-address 10.5.5.0
   ```

2. (Optional) Configure the maximum lifetime that the home agent accepts in any registration request from a mobile node.

   ```
   [edit services mobile-ip home-agent virtual-network]
   user@host# set home-agent-address 10.5.5.0 registration-lifetime 100
   ```

3. (Optional) Configure a timestamp tolerance for registration replay protection.

   ```
   [edit services mobile-ip home-agent virtual-network]
   user@host# set home-agent-address 10.5.5.0 timestamp-tolerance 200
   ```


Copyright © 2013, Juniper Networks, Inc.

4. Configure whether the home agent can revoke a mobile node's registration to deactivate the node.

```
[edit services mobile-ip home-agent virtual-network]
user@host# set home-agent-address 10.5.5.0 revocation-required
```

5. Specify the interfaces on which the home agent accepts registration requests.

```
[edit services mobile-ip home-agent]
user@host# set enable-service ge-0/0/1.0
user@host# set enable-service ge-0/0/2.0
user@host# set enable-service ge-0/0/3.0
user@host# set enable-service ge-0/0/4.0
```

Related Documentation

- Configuring Mobile IP on page 569

Configuring the Local Authentication Attributes for the Mobile Node

You specify for each mobile node several attributes that enable authentication of registration requests from the node. These attributes include security association context for the peering relationship, the entity type of the node, the encryption algorithm and key used to authenticate the request, and replay protection.

To configure authentication attributes for the mobile node:

1. Configure the peer entity for the security parameter.

```
[edit services mobile-ip]
user@host# set peer ip-address 10.4.2.20 spi 500 entity-type mobility-agent
```

2. Configure the algorithm used for authenticating Mobile IP messages. By default, the hmac-md5 algorithm is used.

```
[edit services mobile-ip]
user@host# set peer ip-address 10.4.2.20 spi 500 algorithm md5
```

3. Configure the authentication key for the security association, in either HEX or ASCII format.

```
[edit services mobile-ip]
user@host# set peer ip-address 10.4.2.20 spi 500 key ascii xf125j9m
```

4. Configure a timestamp tolerance for registration replay protection or specify that the timestamp tolerance be taken from the value configured on the home agent.

```
[edit services mobile-ip]
user@host# set peer ip-address 10.4.2.20 spi 500 replay-method timestamp tolerance
  250
```

Related Documentation

- Configuring Mobile IP on page 569

Configuring Accounting for Mobile IP Subscribers

You can configure time-based accounting to track the subscriber sessions of Mobile IP subscribers.

To configure Mobile IP accounting:

1. Configure the IP address for the RADIUS accounting server.

   ```
   [edit access profile mip-win4]
   user@host# set radius accounting-server 192.168.20.5
   ```

2. Specify RADIUS as the accounting method for Mobile IP subscribers.

   ```
   [edit access profile mip-win4 accounting]
   user@host# set order radius
   ```

3. Specify time-based accounting for the access profile used for the subscriber.

   ```
   [edit access profile mip-win4 accounting]
   user@host# set statistics time
   ```

Related Documentation

- Configuring Mobile IP on page 569
- Specifying the Authentication and Accounting Methods for Subscriber Access on page 24
- Configuring Per-Subscriber Session Accounting on page 29
- Configuring RADIUS Server Parameters for Subscriber Access on page 36

Configuring Dynamic Home Assignment for the Mobile Node

The mobile node can request that the home agent dynamically assign an IP address for the home agent. The mobile node uses this address for the home agent in all subsequent registration requests until the registration expires or the mobile node is rebooted.

To configure the IP address to be used by the mobile node for the home agent:

- Configure the IP address for the specified mobile node.

  ```
  [edit services mobile–ip]
  user@host# set dynamic-home-assignment home-agent nai bws@example.com
    home-agent 192.168.4.5
  ```

Related Documentation

- Configuring Mobile IP on page 569

Copyright © 2013, Juniper Networks, Inc.

Configuring the Access Type for Mobile IP

You can configure the Mobile IP home agent to operate in a Worldwide Interoperability for Microwave Access (WiMAX) home connectivity services network (H-CSN). This configuration enables the home agent to receive, process, and send WiMAX VSAs for subscriber authentication and registration. By default, Mobile IP cannot process the WiMAX VSAs. For operation in non-WiMAX environments, you can return it to this mode by configuring the **generic** access type.

NOTE: **The Mobile IP configuration for WiMAX requires that AAA be used for the authentication method. For that reason, WiMAX is available only in the default router context.**

To configure the access type, do one of the following:

- Configure generic operation.

  [edit services mobile-ip]
  user@host# **set** access-type generic

- Configure WiMAX operation.

  [edit services mobile-ip]
  user@host# **set** access-type wimax

Related Documentation

- Configuring Mobile IP on page 569

Tracing Mobile IP Operations for Subscriber Access

The Junos OS trace feature tracks Mobile IP operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

Trace-related configurations are independent for each logical system and routing instance in which Mobile IP is configured. Mobile IP can generate two types of log messages:

- Trace messages common to all logical systems and routing instances in which Mobile IP is configured. Examples of this global message type are the messages generated by Mobile IP during initialization after it starts up. These trace messages are stored in the default trace file, **/var/log/mipd**. You cannot configure Mobile IP to save global messages in a different file. Mobile IP traces global messages by default.
- Trace messages specific to a logical system or routing instance in which Mobile IP is configured. An example of this message type is the message generated by Mobile IP when it receives a registration request. These trace messages are stored in the trace file configured for that logical system or routing instance. These messages cannot be saved in **/var/log/mipd**.

Copyright © 2013, Juniper Networks, Inc.

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 **mipd** for global tracing. You can specify a different filename, but you cannot change the directory in which trace files are located. Logical system and routing instance messages are logged in a file that you must configure separately from **mipd** in the **/var/log** directory.
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 *Junos OS System Log Messages Reference*.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure Mobile IP tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the Mobile IP Trace Log Filename" on page 575.
2. (Optional) Configure the number and size of trace logs.

   See "Configuring the Number and Size of Mobile IP Log Files" on page 575.
3. (Optional) Configure user access to trace logs.

   See "Configuring Access to the Mobile IP Log File" on page 576.
4. (Optional) Configure a regular expression to filter the information to be included in the trace log.

   See "Configuring a Regular Expression for Mobile IP Messages to Be Logged" on page 576.
5. (Optional) Configure flags to specify which events are logged.

   See "Configuring the Mobile IP Tracing Flags" on page 576.
6. (Optional) Configure a severity level for messages to specify which event messages are logged.

   See "Configuring the Severity Level to Filter Which Mobile IP Messages Are Logged" on page 577.

Copyright © 2013, Juniper Networks, Inc.

Configuring the Mobile IP Trace Log Filename

Global messages common to all Mobile IP logical systems and routing instances are recorded only in **/var/log/mipd**. Mobile IP automatically creates this file if it is not present when Mobile IP starts. You cannot configure global messages to be recorded in any other file.

You must specify a different name with the **file** option for messages that are specific to a logical system or routing instance in which Mobile IP is configured. Ensure that filenames are unique for each logical system or routing instance in which Mobile IP is configured. If you do not configure a trace filename for a logical system or routing instance, then nothing is traced for that entity.

To configure the filename for Mobile IP tracing operations for a logical system or routing instance:

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

  [edit logical-systems lr1 services mobile-ip traceoptions]
  user@host# **set file mip-lr1_1**

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 573

Configuring the Number and Size of Mobile IP 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. (Mobile IP supports the **files** and **size** options for the **traceoptions** statement.)

  [edit services mobile-ip traceoptions]
  user@host# **set file mip_1 _logfile_1 files 20 size 2097152**

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 573

Configuring Access to the Mobile IP Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

To specify that all users can read the log file:

- Configure the log file to be world-readable.

```
[edit services mobile-ip traceoptions]
user@host# set file mip_1 _logfile_1 world-readable
```

To explicitly set the default behavior, only the user who configured tracing can read the log file:

- Configure the log file to be no-world-readable.

```
[edit services mobile-ip traceoptions]
user@host# set file mip_1 _logfile_1 no-world-readable
```

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 573

Configuring a Regular Expression for Mobile IP 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 that will be matched.

To configure regular expressions to be matched:

- Configure the regular expression.

```
[edit services mobile-ip traceoptions]
user@host# set file mip_1 _logfile_1 match regex
```

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 573

Configuring the Mobile IP Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

```
[edit services mobile-ip traceoptions]
```

 Copyright © 2013, Juniper Networks, Inc.

```
user@host# set flag home-agent
```

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 573

Configuring the Severity Level to Filter Which Mobile IP Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all** or **verbose**. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages . By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

  ```
  [edit services mobile-ip traceoptions]
  user@host# set level severity
  ```

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 573

Copyright © 2013, Juniper Networks, Inc.

PART 9

ATM for Subscriber Access

- ATM for Subscriber Access Overview on page 581
- Configuring ATM for Subscriber Access on page 587
- ATM for Subscriber Access Examples on page 593

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 30

ATM for Subscriber Access Overview

- ATM for Subscriber Access Overview on page 581
- ATM for Subscriber Access Encapsulation Types Overview on page 585

ATM for Subscriber Access Overview

By using the ATM Modular Interface Card (MIC) with small form-factor pluggable transceiver (SFP) and a supported Modular Port Concentrator (MPC), you can configure the MX Series router to support configurations that enable subscribers to access the router over an ATM network using ATM Adaptation Layer 5 (AAL5) permanent virtual connections (PVCs). Using these configurations enables the delivery of subscriber-based services, such as class of service (CoS) and firewall filters, for subscribers accessing the router over an ATM network.

This overview describes the following topics about configuring ATM interfaces for subscriber access:

- Supported Configurations for ATM Subscriber Access on page 581
- PPP-over-Ethernet-over-ATM Configurations on page 582
- Routed IP-over-ATM Configurations on page 582
- Bridged IP-over-Ethernet-over-ATM Configurations on page 583
- PPP-over-ATM Configurations on page 583
- Configuration and Encapsulation Types for ATM Subscriber Access on page 584

Supported Configurations for ATM Subscriber Access

On MX Series routers with MPC/MIC interfaces that use the ATM MIC with SFP (Model Number MIC-3D-8OC3-2OC12-ATM), you can create the following configurations to enable subscribers to access the router over an ATM network using ATM Adaptation Layer 5 (AAL5) permanent virtual connections (PVCs):

- PPP-over-Ethernet-over-ATM
- Routed IP-over-ATM
- Bridged IP-over-Ethernet-over-ATM
- PPP-over-ATM

Copyright © 2013, Juniper Networks, Inc.

The following sections briefly describe each supported ATM subscriber access configuration.

PPP-over-Ethernet-over-ATM Configurations

PPP-over-Ethernet-over-ATM (PPPoE-over-ATM) configurations support both statically created and dynamically created PPPoE (**pp0**) logical subscriber interfaces over static ATM underlying interfaces. Most PPPoE and subscriber services features supported on terminated connections and tunneled (L2TP access concentrator, or LAC) connections are also supported for access to an MX Series router over an ATM network.

PPPoE-over-ATM configurations require static configuration of the underlying ATM physical interface and ATM logical interface. You can configure the PPPoE (**pp0**) subscriber interface either dynamically, by means of a dynamic profile, or statically at the **[edit interfaces pp0 unit *logical-unit-number*]** hierarchy level.

For PPPoE-over-ATM configurations on an MX Series router, you must configure the ATM underlying interface with PPPoE-over-ATM logical link control (LLC) encapsulation. To do so, include the **encapsulation ppp-over-ether-over-atm-llc** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Using dynamic PPPoE-over-ATM configurations for ATM subscriber access enables you to configure an MX Series router to dynamically create PPPoE logical subscriber interfaces over static ATM underlying interfaces only when needed; that is, when a subscriber logs in on the associated underlying interface. Dynamic PPPoE over static ATM configurations are *not* supported on M Series routers and T Series routers.

Optionally, you can dynamically or statically apply subscriber services such as class of service (CoS) and firewall filters to the PPPoE (**pp0**) subscriber interface. For PPPoE-over-ATM configurations that create a dynamic PPPoE subscriber interface, you can configure CoS attributes and firewall filters in the dynamic profile that defines the **pp0** subscriber interface. For PPPoE-over-ATM configurations that create a static PPPoE subscriber interface, you can statically configure CoS attributes and firewall filters as you would for any static interface configured on an MX Series router.

Routed IP-over-ATM Configurations

Routed IP-over-ATM (IPoA) configurations support statically created IPv4 and IPv6 logical subscriber interfaces over static ATM underlying interfaces. IPoA configurations are typically used to implement business digital subscriber line (DSL) connections that do not require connection negotiation for address assignment.

IPoA configurations on MX Series routers require static configuration of the ATM underlying interface, IPv4 interface, IPv6 interface, CoS attributes, and firewall fiters. Dynamic configuration of these components is not supported.

To configure IPoA subscriber access on MX Series routers, you must configure either of the following encapsulation types on the ATM underlying interface:

- For IPoA encapsulation with logical link control (LLC), configure ATM subnetwork attachment point (SNAP) encapsulation by including the **encapsulation atm-snap**

 Copyright © 2013, Juniper Networks, Inc.

statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

- For IPoA encapsulation with virtual circuit (VC) multiplexing, configure ATM VC multiplex encapsulation by including the **encapsulation atm-vc-mux** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Optionally, you can statically configure subscriber services such as CoS and firewall filters and apply them to the IPv4 or IPv6 interface; you *cannot* use a dynamic profile for this purpose.

Bridged IP-over-Ethernet-over-ATM Configurations

Bridged IP-over-Ethernet-over-ATM (IPoE-over-ATM) configurations support statically created IPv4 and IPv6 logical subscriber interfaces over static ATM underlying interfaces. Like IPoA configurations, IPoE-over-ATM configurations are typically used in topologies that do not require connection negotiation for address assignment.

IPoE-over-ATM configurations on MX Series routers require static configuration of the ATM underlying interface, IP interface, CoS attributes, and firewall fiters. Dynamic configuration of these components is not supported.

For bridged IP-over-Ethernet-over-ATM configurations on an MX Series router, you must configure the ATM underlying interface with Ethernet-over-ATM LLC encapsulation. To do so, include the **encapsulation ether-over-atm-llc** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Optionally, you can statically configure subscriber services such as class of service (CoS) and firewall filters and apply them to the IPv4 or IPv6 interface; you *cannot* use a dynamic profile for this purpose.

PPP-over-ATM Configurations

PPP-over-ATM (PPPoA) configurations support statically created PPP logical subscriber interfaces over static ATM underlying interfaces. Most features supported for PPPoE configurations are also supported for PPP access to an MX Series router over an ATM network.

PPPoA configurations on MX Series routers require static configuration of the ATM underlying interface and PPP subscriber interface.

To configure PPPoA subscriber access on MX Series routers, you must configure either of the following encapsulation types on each PPP logical subscriber interface:

- For PPPoA encapsulation with logical link control (LLC), configure PPP-over-AAL5 LLC encapsulation by including the **encapsulation atm-ppp-llc** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.
- For PPPoA encapsulation with virtual circuit (VC) multiplexing, configure PPP-over-AAL5 multiplex encapsulation by including the **encapsulation atm-ppp-vc-mux** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Optionally, you can use dynamic profiles to dynamically or statically apply subscriber services, such as CoS attributes and firewall filters, to the static PPP subscriber interface. Configuring CoS and firewall filters in this manner enables you to efficiently and economically provide these services to PPP subscribers accessing the router over an ATM network.

Configuration and Encapsulation Types for ATM Subscriber Access

You use the same basic statements, commands, and procedures to create, verify, and manage PPPoE-over-ATM, IPoA, IPoE-over-ATM, and PPPoA configurations as the statements, commands, and procedures you use for static configurations on M Series routers and T Series routers, and for dynamic PPPoE configurations on MX Series routers.

A critical element of configuring ATM subscriber access is ensuring that you specify the correct encapsulation type for the ATM logical interface. The encapsulation type you use depends on the supported configuration and, for IPoA and PPPoA configurations, whether you want to configure an encapsulation type that uses logical link control (LLC) or virtual circuit (VC) multiplexing.

Related Documentation

- ATM for Subscriber Access Encapsulation Types Overview on page 585
- Configuring ATM for Subscriber Access on page 587
- Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM on page 593
- Example: Configuring a Static PPPoE Subscriber Interface over ATM on page 602
- Example: Configuring a Static Subscriber Interface for IP Access over ATM on page 610
- Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM on page 616
- Example: Configuring a Static PPP Subscriber Interface over ATM on page 623

Copyright © 2013, Juniper Networks, Inc.

ATM for Subscriber Access Encapsulation Types Overview

To enable subscriber access to an MX Series router over an ATM network, you can create any of the following configurations on Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces that use the ATM MIC with SFP:

- PPP-over-Ethernet-over-ATM (PPPoE-over ATM) with a dynamic or static PPPoE (**pp0**) subscriber interface over a static ATM underlying interface
- Routed IP-over-ATM (IPoA) with a static IPv4 or IPv6 subscriber interface over a static ATM underlying interface
- Bridged IP-over-Ethernet-over-ATM (IPoE-over-ATM) with a static IPv4 or IPv6 subscriber interface over a static ATM underlying interface
- PPP-over-ATM (PPPoA) with a static PPP subscriber interface over a static ATM underlying interface

As part of the configuration procedure, you must specify the appropriate encapsulation type for your configuration on the ATM logical interface.

Table 58 on page 585 lists and describes the encapsulation type you must specify as part of the **encapsulation** statement when you configure the ATM logical interface for each supported configuration.

Table 58: Encapsulation Types for Supported ATM Subscriber Access Configurations

ATM Subscriber Access Configuration	Encapsulation Type	Description
PPPoE-over-ATM with dynamic **pp0** subscriber interface	**ppp-over-ether-over-atm-llc**	PPPoE-over-ATM encapsulation with logical link control (LLC)
PPPoE-over-ATM with static **pp0** subscriber interface	**ppp-over-ether-over-atm-llc**	PPPoE-over-ATM encapsulation with LLC
IP-over-ATM (IPoA)	**atm-snap**	ATM subnetwork attachment point (SNAP) encapsulation for IPoA with LLC
	atm-vc-mux	ATM VC multiplex encapsulation for IPoA with virtual circuit (VC) multiplexing
IP-over-Ethernet-over-ATM (IPoE-over-ATM)	**ether-over-atm-llc**	Ethernet-over-ATM encapsulation with LLC

Copyright © 2013, Juniper Networks, Inc.

Table 58: Encapsulation Types for Supported ATM Subscriber Access Configurations *(continued)*

ATM Subscriber Access Configuration	Encapsulation Type	Description
PPP-over-ATM (PPPoA)	**atm-ppp-llc** (for PPPoA with logical link control)	PPP-over-AAL5 encapsulation with LLC
	atm-ppp-vc-mux (for PPPoA with virtual circuit multiplexing)	PPP-over-AAL5 encapsulation with VC multiplexing

Related Documentation

- ATM for Subscriber Access Overview on page 581
- Configuring ATM for Subscriber Access on page 587
- Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM on page 593
- Example: Configuring a Static PPPoE Subscriber Interface over ATM on page 602
- Example: Configuring a Static Subscriber Interface for IP Access over ATM on page 610
- Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM on page 616
- Example: Configuring a Static PPP Subscriber Interface over ATM on page 623

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 31

Configuring ATM for Subscriber Access

- Configuring ATM for Subscriber Access on page 587
- Guidelines for Configuring ATM for Subscriber Access on page 589
- Verifying and Managing ATM Configurations for Subscriber Access on page 590

Configuring ATM for Subscriber Access

On MX Series routers with MPC/MIC interfaces that use the ATM MIC with SFP, you can create the following configurations to enable subscribers to access the router over an ATM network using ATM Adaptation Layer 5 (AAL5) permanent virtual connections (PVCs):

- PPP-over-Ethernet-over-ATM (PPPoE-over ATM) with a dynamic PPPoE (**pp0**) subscriber interface over a static ATM underlying interface
- PPP-over-Ethernet-over-ATM (PPPoE-over ATM) with a static PPPoE (**pp0**) subscriber interface over a static ATM underlying interface
- Routed IP-over-ATM (IPoA) with a static IPv4 or IPv6 subscriber interface over a static ATM underlying interface
- Bridged IP-over-Ethernet-over-ATM with a static IPv4 or IPv6 subscriber interface over a static ATM underlying interface
- PPP-over-ATM (PPPoA) with a static PPP subscriber interface over a static ATM underlying interface

Before you begin:

1. Make sure the MX Series router you are using has Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces and an ATM MIC with SFP (Model Number MIC-3D-8OC3-2OC12-ATM) installed and operational.
   - For information about compatible MPCs for the ATM MIC with SFP, see the *MX Series Interface Module Reference*.
   - For information about installing MPCs and MICs in an MX Series router, see the *Hardware Guide* for your MX Series router model.
2. Make sure you understand how to configure and use static ATM interfaces.

   See *ATM Interfaces Overview*.

3. If your configuration includes dynamic profiles for PPPoE, class of service (CoS) attributes, or standard firewall filters, make sure you understand how to configure these attributes and apply them to the subscriber interface.
   - For PPPoE dynamic profiles, see "Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles" on page 899
   - For CoS configuration, see "Configuring Traffic Scheduling and Shaping for Subscriber Access" on page 971
   - For standard firewall flter configuration, see *Guidelines for Configuring Firewall Filters* and *Guidelines for Applying Firewall Filters*

To configure ATM for subscriber access on an MX Series router:

1. For a PPPoE-over-ATM configuration with a dynamic PPPoE (**pp0**) subscriber interface, create a dynamic profile that defines the **pp0** subscriber interface.

   See "Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM" on page 593.
2. Configure one or more virtual path identifiers (VPIs) on the ATM physical interface.
3. Configure the ATM logical subscriber interface.
   a. Configure the appropriate encapsulation type for your configuration.

      See "ATM for Subscriber Access Encapsulation Types Overview" on page 585.
   b. Configure a virtual circuit identifier (VCI) for each VPI configured on the ATM logical interface.
   c. Configure other interface-specific properties as needed for your configuration.

      See "Guidelines for Configuring ATM for Subscriber Access" on page 589.
4. For static PPPoE-over-ATM configurations, define the static PPPoE (**pp0**) subscriber interface at the **[edit interfaces pp0 unit *logical-unit-number*]** hierarchy level.

   See "Example: Configuring a Static PPPoE Subscriber Interface over ATM" on page 602.
5. (Optional) Verify the configuration for ATM subscriber access.

   See "Verifying and Managing ATM Configurations for Subscriber Access" on page 590.

Related Documentation

- ATM for Subscriber Access Overview on page 581
- ATM for Subscriber Access Encapsulation Types Overview on page 585
- Guidelines for Configuring ATM for Subscriber Access on page 589
- Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM on page 593
- Example: Configuring a Static PPPoE Subscriber Interface over ATM on page 602
- Example: Configuring a Static Subscriber Interface for IP Access over ATM on page 610
- Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM on page 616

Copyright © 2013, Juniper Networks, Inc.

- Example: Configuring a Static PPP Subscriber Interface over ATM on page 623
- *ATM Interfaces Overview*

Guidelines for Configuring ATM for Subscriber Access

The following guidelines apply when you configure PPP-over-Ethernet-over-ATM (PPPoE-over-ATM), IP-over-ATM (IPoA), IP-over-Ethernet-over-ATM (IPoE-over-ATM), or PPP-over-ATM (PPPoA) configurations for ATM subscriber access. You can create these configurations on MX Series routers with Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces that use the ATM MIC with SFP.

For all supported ATM subscriber access configurations:

- Make sure you specify the correct encapsulation type on the ATM logical interface for your configuration, as described in "ATM for Subscriber Access Encapsulation Types Overview" on page 585.

For PPPoE-over-ATM configurations:

- For dynamic or static PPPoE-over-ATM configurations, specify PPPoE-specific options at the **[edit interfaces *interface-name* unit *logical-unit-number* family pppoe]** hierarchy level. Specifying PPPoE-specific options at the **[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]** hierarchy level is not supported for these configurations.
- For dynamic or static PPPoE-over-ATM configurations, you must configure the router to act as a PPPoE server (also known as a *remote access concentrator*). Configuring the router to act as a PPPoE client is not supported in these configurations.
- For dynamic PPPoE-over-ATM configurations, issue the **dynamic-profile *profile-name*** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family pppoe]** hierarchy level to associate the ATM logical interface with the dynamic profile that defines the PPPoE subscriber interface.

For static IPoA and IPoE-over-ATM configurations:

- Specify interface-specific options at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level (for IPv4) or at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet6]** hierarchy level (for IPv6).

For static PPPoA configurations:

- Specify PPP-specific options at the **[edit interfaces *interface-name* unit *logical-unit-number* ppp-options]** hierarchy level.

Related Documentation

- ATM for Subscriber Access Overview on page 581
- ATM for Subscriber Access Encapsulation Types Overview on page 585
- Configuring ATM for Subscriber Access on page 587
- Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM on page 593

Copyright © 2013, Juniper Networks, Inc.

- Example: Configuring a Static PPPoE Subscriber Interface over ATM on page 602
- Example: Configuring a Static Subscriber Interface for IP Access over ATM on page 610
- Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM on page 616
- Example: Configuring a Static PPP Subscriber Interface over ATM on page 623

Verifying and Managing ATM Configurations for Subscriber Access

Purpose View information about the static or dynamic subscriber interfaces configured over a static ATM underlying interface on an MX Series router with MPC/MIC interfaces and an ATM MIC with SFP.

Action

- To display information about the ATM physical interface to ensure that it is properly configured for use with ATM PVCs:

  `user@host>` **show interfaces at-*fpc/pic/port***

- To display information about the ATM logical interface to ensure that it is properly configured as a dynamic or static subscriber interface:

  `user@host>` **show interfaces at-*fpc/pic/port.logical-unit-number***

- To display information about all static PPPoE (**pp0**) subscriber interfaces for static PPPoE-over-ATM configurations:

  `user@host>` **show interfaces pp0**

- To display information about a specified static PPPoE (**pp0**) subscriber interface for static PPPoE-over-ATM configurations:

  `user@host>` **show interfaces pp0.*logical-unit-number***

- To display detailed information about the PPPoE underlying interface for dynamic or static PPPoE-over-ATM configurations:

  `user@host>` **show pppoe underlying-interfaces at-*fpc/pic/port.logical-unit-number* detail**

- To display extensive information, including packet statistics and lockout time settings, about the PPPoE underlying interface for dynamic or static PPPoE-over-ATM configurations:

  `user@host>` **show pppoe underlying-interfaces at-*fpc/pic/port.logical-unit-number* extensive**

- To display extensive information about the active ATM subscriber with the specified ATM virtual path identifier (VPI) and ATM virtual circuit identifier (VCI):

  `user@host>` **show subscribers vpi *vpi-identifier* vci *vci-identifier* extensive**

Related Documentation

- Configuring ATM for Subscriber Access on page 587
- Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM on page 593
- Example: Configuring a Static PPPoE Subscriber Interface over ATM on page 602
- Example: Configuring a Static Subscriber Interface for IP Access over ATM on page 610
- Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM on page 616

 Copyright © 2013, Juniper Networks, Inc.

- Example: Configuring a Static PPP Subscriber Interface over ATM on page 623

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 32

ATM for Subscriber Access Examples

- Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM on page 593
- Example: Configuring a Static PPPoE Subscriber Interface over ATM on page 602
- Example: Configuring a Static Subscriber Interface for IP Access over ATM on page 610
- Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM on page 616
- Example: Configuring a Static PPP Subscriber Interface over ATM on page 623

Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM

This example illustrates a Point-to-Point Protocol over Ethernet (PPPoE) over ATM configuration that creates a dynamic PPPoE (**pp0**) subscriber interface over a static ATM underlying interface on an MX Series router. The router must have Module Port Concentrator/Modular Interface Card (MPC/MIC) interfaces that use an ATM MIC with small form-factor pluggable transceiver (SFP).

NOTE: You can also configure a *static* PPPoE interface over a static ATM underlying interface on an MX Series router with an ATM MIC with SFP installed. For information, see "Example: Configuring a Static PPPoE Subscriber Interface over ATM" on page 602.

- Requirements on page 593
- Overview on page 594
- Configuration on page 595
- Verification on page 600

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.2
- MX Series 3D Universal Edge Router
- ATM MIC with SFP (Model Number MIC-3D-8OC3-2OC12-ATM) and compatible MPC1 or MPC2

Before you begin:

1. Make sure the MX Series router you are using has an ATM MIC with SFP installed and operational.
   - For information about compatible MPCs for the ATM MIC with SFP, see the *MX Series Interface Module Reference*.
   - For information about installing MPCs and MICs in an MX Series router, see the *Hardware Guide* for your MX Series router model.
2. Make sure you understand how to configure and use static ATM interfaces.

   See *ATM Interfaces Overview*.
3. Make sure you understand how to configure and use dynamic PPPoE subscriber interfaces.
   - For overview information, see "Subscriber Interfaces and PPPoE Overview" on page 883
   - For configuration instructions, see "Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles" on page 899

Overview

By using the ATM MIC with SFP and a supported MPC, you can configure an MX Series router to support dynamic PPPoE subscriber access over an ATM network. PPPoE-over-ATM configurations on MX Series routers consist of one or more dynamically created PPPoE (**pp0**) subscriber interfaces over a static ATM underlying interface. Most PPPoE and subscriber services features supported on terminated connections and tunneled (L2TP access concentrator, or LAC) connections are also supported for PPPoE-over-ATM connections on an MX Series router.

Optionally, you can dynamically apply subscriber services such as class of service (CoS) and firewall filters to the PPPoE subscriber interface by configuring these services in the dynamic profile that creates the **pp0** subscriber interface. In this example, the PPPoE dynamic profile (pppoe-profile) applies CoS traffic shaping parameters to the dynamic **pp0** subscriber interface. Configuring CoS and firewall filters in this manner enables you to efficiently and economically provide these services to PPPoE subscribers accessing the router over an ATM network using ATM Adaptation Layer 5 (AAL5) permanent virtual connections (PVCs).

This example includes the following basic steps to configure dynamic PPPoE-over-ATM subscriber access on an MX Series router:

1. Create a PPPoE dynamic profile named pppoe-profile for the **pp0** subscriber interface that includes all of the following:
   - The logical unit number, represented by the **$junos-interface-unit** predefined dynamic variable
   - The name of the underlying ATM interface, represented by the **$junos-underlying-interface** predefined dynamic variable

 Copyright © 2013, Juniper Networks, Inc.

- The **server** statement, which configures the router to act as a PPPoE server

  NOTE: **Configuring the router to act as a PPPoE client is not supported.**

- The unnumbered address (lo0.0) for the IPv4 (**inet**) protocol family
- CoS traffic shaping parameters

2. Statically configure the ATM physical interface at-1/0/0 with virtual path identifier (VPI) 3.
3. Statically configure logical unit 2 on the ATM physical interface (at-1/0/0.2) with at least the following properties:
   - PPPoE-over-ATM logical link control (LLC) encapsulation (**ppp-over-ether-over-atm-llc**)
   - Virtual circuit identifier (VCI) 2 on VPI 3. The combination of VPIs and VCIs provisions the ATM AAL5 PVC for access over the ATM network.
   - PPPoE-specific options at the **[edit interfaces *interface-name* unit *logical-unit-number* family pppoe]** hierarchy level, including at least the name of the associated PPPoE dynamic profile (pppoe-profile) that creates the pp0 dynamic subscriber interface

     In dynamic PPPoE-over-ATM configurations, each **pp0** interface defined in the dynamic profile corresponds to a dynamic PPPoE subscriber interface.

     NOTE: **For dynamic or static PPPoE-over-ATM configurations on MX Series routers, You must specify PPPoE-specific options in the family pppoe stanza at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. Specifying PPPoE-specific options in the pppoe-underlying-options stanza at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level is not supported for these configurations.**

Configuration

To configure a dynamic PPPoE subscriber interface over an underlying ATM interface, perform these tasks:

- Configuring the PPPoE Dynamic Profile on page 596
- Configuring the ATM Physical Interface on page 598
- Configuring the Dynamic PPPoE Subscriber Interface on Logical Unit 2 on page 599

CLI Quick Configuration

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

```
# PPPoE Dynamic Profile
```

```
set dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit" ppp-options
  chap
set dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit"
  pppoe-options underlying-interface "$junos-underlying-interface"
set dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit"
  pppoe-options server
set dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit"
  no-keepalives
set dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit" family
  inet unnumbered-address lo0.0
set dynamic-profiles pppoe-profile class-of-service traffic-control-profiles tcp-test
  shaping-rate 10m
set dynamic-profiles pppoe-profile class-of-service interfaces pp0 unit
  "$junos-interface-unit" output-traffic-control-profile tcp-test
#
# ATM Physical Interface
set interfaces at-1/0/0 atm-options vpi 3
#
# Logical Unit 2
set interfaces at-1/0/0 atm-options vpi 3
set interfaces at-1/0/0 unit 2 encapsulation ppp-over-ether-over-atm-llc
set interfaces at-1/0/0 unit 2 vci 3.2
set interfaces at-1/0/0 unit 2 family pppoe access-concentrator ac-pppoeoa
set interfaces at-1/0/0 unit 2 family pppoe duplicate-protection
set interfaces at-1/0/0 unit 2 family pppoe dynamic-profile pppoe-profile
set interfaces at-1/0/0 unit 2 family pppoe max-sessions 3
set interfaces at-1/0/0 unit 2 family pppoe short-cycle-protection
```

Configuring the PPPoE Dynamic Profile

Step-by-Step Procedure

To configure the PPPoE dynamic profile for the **pp0** subscriber interface:

1. Name the dynamic profile.

   ```
   [edit]
   user@host# edit dynamic-profiles pppoe-profile
   ```

2. Specify that you want to configure the **pp0** (PPPoE) interface.

   ```
   [edit dynamic-profiles pppoe-profile]
   user@host# edit interfaces pp0
   ```

3. Specify that you want to configure the logical unit represented by the **$junos-interface-unit** predefined variable.

   ```
   [edit dynamic-profiles pppoe-profile interfaces pp0]
   user@host# edit unit $junos-interface-unit
   ```

   The **$junos-interface-unit** variable is dynamically replaced with the actual unit number supplied by the network when the subscriber logs in.

4. Configure PPPoE-specific options for the **pp0** interface.

   a. Configure the ATM underlying interface represented by the **$junos-underlying-interface** predefined variable.

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


Copyright © 2013, Juniper Networks, Inc.

The **$junos-underlying-interface** variable is dynamically replaced with the actual name of the underlying interface supplied by the network when the subscriber logs in.

b. Configure the router to act as a PPPoE server, also known as a remote access concentrator.

```
[edit dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options server
```

5. Configure Challenge Handshake Authentication Protocol (CHAP) authentication for the **pp0** interface.

```
[edit dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set ppp-options chap
```

6. Disable sending keepalive messages on the interface.

```
[edit dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set no-keepalives
```

7. Configure the protocol family for the **pp0** interface.

a. Specify that you want to configure the IPv4 (**inet**) protocol family.

```
[edit dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# edit family inet
```

b. Configure the unnumbered address for the protocol family.

```
[edit dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit"
  family inet]
user@host# set unnumbered-address lo0.0
user@host# up 4
```

8. Configure CoS traffic shaping parameters in the dynamic profile for the pp0 subscriber interface.

a. Specify that you want to configure CoS traffic shaping parameters.

```
[edit dynamic-profiles pppoe-profile]
user@host# edit class-of-service
```

b. Create a traffic control profile.

```
[edit dynamic-profiles pppoe-profile class-of-service]
user@host# edit traffic-control-profiles tcp-test
```

c. Configure the traffic shaping rate.

```
[edit dynamic-profiles pppoe-profile class-of-service traffic-control-profiles
  tcp-test]
user@host# set shaping-rate 10m
user@host# up 2
```

d. Apply the traffic shaping parameters to the pp0 dynamic subscriber interface.

```
[edit dynamic-profiles pppoe-profile class-of-service]
user@host# edit interfaces pp0 unit $junos-interface-unit
```

e. Apply the output traffic scheduling and shaping profile to the interface.

```
[edit dynamic-profiles pppoe-profile class-of-service interfaces pp0 unit
  "$junos-interface-unit]
user@host# set output-traffic-control-profile tcp-test
```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the PPPoE dynamic profile configuration by issuing the **show dynamic-profiles pppoe-profile** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show dynamic-profiles pppoe-profile
interfaces {
  pp0 {
    unit "$junos-interface-unit" {
      ppp-options {
        chap;
      }
      pppoe-options {
        underlying-interface "$junos-underlying-interface";
        server;
      }
      no-keepalives;
      family inet {
        unnumbered-address lo0.0;
      }
    }
  }
}
class-of-service {
  traffic-control-profiles {
    tcp-test {
      shaping-rate 10m;
    }
  }
  interfaces {
    pp0 {
      unit "$junos-interface-unit" {
        output-traffic-control-profile tcp-test;
      }
    }
  }
}
```

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

Configuring the ATM Physical Interface

Step-by-Step Procedure To configure the ATM physical interface:

1. Specify that you want to configure ATM-specific options on the physical interface.

   ```
   [edit interfaces at-1/0/0]
   user@host# edit atm-options
   ```

2. Configure one or more VPIs on the physical interface.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces at-1/0/0 atm-options]
user@host# set vpi 3
```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the ATM physical interface configuration by issuing the **show interfaces at-1/0/0** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/0
atm-options {
   vpi 3;
}
```

If you are done configuring the ATM physical interface, enter **commit** from configuration mode.

Configuring the Dynamic PPPoE Subscriber Interface on Logical Unit 2

Step-by-Step Procedure To configure the dynamic PPPoE subscriber interface on logical unit 2:

1. Configure PPPoE-over-ATM LLC encapsulation on the interface.

   ```
   [edit interfaces at-1/0/0 unit 2]
   user@host# set encapsulation ppp-over-ether-over-atm-llc
   ```

2. Configure the VCI for the logical interface.

   ```
   [edit interfaces at-1/0/0 unit 2]
   user@host# set vci 3.2
   ```

   This statement configures VCI 2 on VPI 3.

3. Specify that you want to configure the PPPoE protocol family.

   ```
   [edit interfaces at-1/0/0 unit 2]
   user@host# edit family pppoe
   ```

4. Associate the interface with the dynamic profile that creates the dynamic PPPoE subscriber interface.

   ```
   [edit interfaces at-1/0/0 unit 2 family pppoe]
   user@host# set dynamic-profile pppoe-profile
   ```

5. Configure additional PPPoE-specific options for the dynamic subscriber interface.

   ```
   [edit interfaces at-1/0/0 unit 2 family pppoe]
   user@host# set max-sessions 3
   user@host# set duplicate-protection
   user@host# set short-cycle-protection
   user@host# set access-concentrator ac-pppoeoa
   ```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the dynamic PPPoE subscriber interface configuration on logical unit 2 by issuing the **show interfaces at-1/0/0.2** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/0.2
encapsulation ppp-over-ether-over-atm-llc;
vci 3.2;
family pppoe {
   access-concentrator ac-pppoeoa;
   duplicate-protection;
   dynamic-profile pppoe-profile;
   max-sessions 3;
   short-cycle-protection;
}
```

If you are done configuring the dynamic PPPoE subscriber interface on logical unit 2, enter **commit** from configuration mode.

Verification

To confirm that the dynamic PPPoE subscriber interface is properly configured on ATM interface at-1/0/0.2, perform the following tasks:

- Verifying the ATM Physical Interface Configuration on page 600
- Verifying the Dynamic PPPoE Subscriber Interface Configuration on Logical Unit 2 on page 601
- Verifying the PPPoE Underlying Interface Configuration on page 601

Verifying the ATM Physical Interface Configuration

Purpose Verify that ATM physical interface at-1/0/0 is properly configured for use with ATM PVCs.

Action From operational mode, issue the **show interfaces at-1/0/0** command.

For brevity, this **show** command output includes only the configuration that is relevant to the at-1/0/0 physical interface. Any other configuration on the system has been replaced with ellipses (...).

```
user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 592
  Link-level type: ATM-PVC, MTU: 2048, Clocking: Internal, SDH mode, Speed: OC3,
 Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 00:1f:12:bc:4a:95
  Last flapped   : 2012-09-17 07:21:19 PDT (08:26:16 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SDH   alarms   : None
  SDH   defects  : None
    VPI 3
      Flags: Active
      Total down time: 0 sec, Last down: Never
  Traffic statistics:
         Input  packets:                    0
```

 Copyright © 2013, Juniper Networks, Inc.

```
        Output packets:                    0
...
```

Meaning **ATM-PVC** in the Link-level Type field indicates that encapsulation for ATM permanent virtual circuits is being used on ATM physical interface at-1/0/0. The **Active** flag for VPI 3 indicates that the virtual path is up and operational.

Verifying the Dynamic PPPoE Subscriber Interface Configuration on Logical Unit 2

Purpose Verify that the dynamic PPPoE subscriber interface is properly configured on logical unit 2 (at-1/0/0.2).

Action From operational mode, issue the **show interfaces at-1/0/0.2** command.

```
user@host> show interfaces at-1/0/0.2
  Logical interface at-1/0/0.2 (Index 350) (SNMP ifIndex 1701)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE-over-ATM-LLC
    Input packets : 0
    Output packets: 0
    Protocol pppoe
      Dynamic Profile: pppoe-profile,
      Service Name Table: None,
      Max Sessions: 3, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: mac-address,
      AC Name: ac-pppoeoa
    VCI 3.2
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Input packets : 0
        Output packets: 0
```

Meaning **PPPoE-over-ATM-LLC** in the Encapsulation field indicates that logical interface at-1/0/0.2 is properly configured for PPPoE-over-ATM LLC encapsulation. **Protocol pppoe** indicates that the PPPoE protocol family has been properly configured on the logical interface. The Dynamic Profile field indicates that dynamic profile **pppoe-profile** creates the dynamic PPPoE subscriber interface. The **Active** flag for VCI 3.2 indicates that VCI 2 on VPI 3 is up and operational.

Verifying the PPPoE Underlying Interface Configuration

Purpose Verify that the underlying interface is properly configured for dynamic PPPoE-over-ATM subscriber access.

Action From operational mode, issue the **show pppoe underlying-interfaces at-1/0/0.2** command.

```
user@host> show pppoe underlying-interfaces at-1/0/0.2 detail
at-1/0/0.2 Index 350
  State: Static, Dynamic Profile: pppoe-profile,
  Max Sessions: 3, Max Sessions VSA Ignore: Off,
  Active Sessions: 0,
  Service Name Table: None,
  Duplicate Protection: On, Short Cycle Protection: mac-address,
  AC Name: ac-pppoeoa,
```

Meaning This command indicates that ATM logical interface at-1/0/0.2 is properly configured as the PPPoE underlying interface. **Static** in the State field indicates that at-1/0/0/.2 is statically configured. The Dynamic Profile field indicates that **pppoe-profile** is the name of the dynamic profile used to create this interface. The remaining fields display information about the PPPoE-specific interface options configured for the PPPoE underlying interface at the **[edit interfaces at-1/0/0 unit 2 family pppoe]** hierarchy level.

Related Documentation

- ATM for Subscriber Access Overview on page 581
- Configuring ATM for Subscriber Access on page 587
- Example: Configuring a Static PPPoE Subscriber Interface over ATM on page 602
- Example: Configuring a Static Subscriber Interface for IP Access over ATM on page 610
- Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM on page 616
- Example: Configuring a Static PPP Subscriber Interface over ATM on page 623

Example: Configuring a Static PPPoE Subscriber Interface over ATM

This example illustrates a Point-to-Point Protocol over Ethernet (PPPoE) over ATM configuration that creates a static PPPoE (**pp0**) subscriber interface over a static ATM underlying interface on an MX Series router. The router must have Module Port Concentrator/Modular Interface Card (MPC/MIC) interfaces that use an ATM MIC with small form-factor pluggable transceiver (SFP).

NOTE: You can also configure a *dynamic* PPPoE interface over a static ATM underlying interface on an MX Series router with an ATM MIC with SFP installed. For information, see "Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM" on page 593.

- Requirements on page 602
- Overview on page 603
- Configuration on page 604
- Verification on page 607

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.2
- MX Series 3D Universal Edge Router
- ATM MIC with SFP (Model Number MIC-3D-8OC3-2OC12-ATM) and compatible MPC1 or MPC2

 Copyright © 2013, Juniper Networks, Inc.

Before you begin:

1. Make sure the MX Series router you are using has an ATM MIC with SFP installed and operational.
   - For information about compatible MPCs for the ATM MIC with SFP, see the *MX Series Interface Module Reference*.
   - For information about installing MPCs and MICs in an MX Series router, see the *Hardware Guide* for your MX Series router model.
2. Make sure you understand how to configure and use static ATM interfaces.

   See *ATM Interfaces Overview*.

Overview

By using the ATM MIC with SFP and a supported MPC, you can configure an MX Series router to support static PPPoE subscriber access over an ATM network using ATM Adaptation Layer 5 (AAL5) permanent virtual connections (PVCs). PPPoE-over-ATM configurations on MX Series routers consist of one or more statically created PPPoE (**pp0**) logical subscriber interfaces over a static ATM underlying interface. Most PPPoE and subscriber services features supported on terminated connections and tunneled (L2TP access concentrator, or LAC) connections are also supported for PPPoE-over-ATM connections on an MX Series router.

This example include the following basic steps to configure static PPPoE-over-ATM subscriber access on an MX Series router:

1. Statically configure ATM physical interface at-1/0/6 with virtual path identifier (VPI) 6.
2. Statically configure logical unit 2 on the ATM physical interface (at-1/0/6.2) with the following properties:
   - PPPoE-over-ATM logical link control (LLC) encapsulation (**ppp-over-ether-over-atm-llc**)
   - Virtual circuit identifier (VCI) 2 on VPI 6. The combination of VPIs and VCIs provisions the ATM AAL5 PVC for access over the ATM network.
   - (Optional) PPPoE-specific options at the **[edit interfaces *interface-name* unit *logical-unit-number* family pppoe]** hierarchy level

   NOTE: For dynamic or static PPPoE-over-ATM configurations on MX Series routers, You must specify PPPoE-specific options in the **family pppoe** stanza at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level. Specifying PPPoE-specific options in the **pppoe-underlying-options** stanza at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level is not supported for these configurations.

3. Statically configure the **pp0** logical subscriber interface (pp0.2) with at least the following properties:
   - The name of the underlying ATM interface (at-1/0/6.2)
   - The **server** statement, which configures the router to act as a PPPoE server
   - The unnumbered address (lo0.0) for the **inet** (IPv4) or **inet6** (IPv6) protocol family

   In static PPPoE-over-ATM configurations, each **pp0** logical interface configured at the **[edit interfaces pp0 unit *logical-unit-number*]** hierarchy level corresponds to a static PPPoE subscriber interface.

Configuration

To configure a static PPPoE subscriber interface over an underlying ATM interface, perform these tasks:

- Configuring the ATM Physical Interface on page 604
- Configuring Encapsulation, VCI, and PPPoE Options on Logical Unit 2 on page 605
- Configuring the Static PPPoE Subscriber Interface on page 606

CLI Quick Configuration

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

```
# ATM Physical Interface
set interfaces at-1/0/6 atm-options vpi 6
#
# Logical Unit 2
set interfaces at-1/0/6 unit 2 encapsulation ppp-over-ether-over-atm-llc
set interfaces at-1/0/6 unit 2 vci 6.2
set interfaces at-1/0/6 unit 2 family pppoe access-concentrator ac-pppoeoa
set interfaces at-1/0/6 unit 2 family pppoe duplicate-protection
set interfaces at-1/0/6 unit 2 family pppoe max-sessions 3
set interfaces at-1/0/6 unit 2 family pppoe max-sessions-vsa-ignore
set interfaces at-1/0/6 unit 2 family pppoe short-cycle-protection lockout-time-min 120
set interfaces at-1/0/6 unit 2 family pppoe short-cycle-protection lockout-time-max 240
#
# Static PPPoE Subscriber Interface
set interfaces pp0 unit 2 ppp-options chap
set interfaces pp0 unit 2 pppoe-options underlying-interface at-1/0/6.2
set interfaces pp0 unit 2 pppoe-options server
set interfaces pp0 unit 2 keepalives interval 10
set interfaces pp0 unit 2 family inet unnumbered-address lo0.0
```

Configuring the ATM Physical Interface

Step-by-Step Procedure

To configure the ATM physical interface:

1. Specify that you want to configure ATM-specific options on the physical interface.

   ```
   [edit interfaces at-1/0/6]
   user@host# edit atm-options
   ```


Copyright © 2013, Juniper Networks, Inc.

2. Configure one or more VPIs on the physical interface.

```
[edit interfaces at-1/0/6 atm-options]
user@host# set vpi 6
```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the ATM physical interface configuration by issuing the **show interfaces at-1/0/6** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/6
atm-options {
   vpi 6;
}
```

If you are done configuring the ATM physical interface, enter **commit** from configuration mode.

Configuring Encapsulation, VCI, and PPPoE Options on Logical Unit 2

Step-by-Step Procedure To configure encapsulation, VCI, and PPPoE options on logical unit 2:

1. Configure PPPoE-over-ATM LLC encapsulation on the interface.

```
[edit interfaces at-1/0/6 unit 2]
user@host# set encapsulation ppp-over-ether-over-atm-llc
```

2. Configure the VCI for the logical interface.

```
[edit interfaces at-1/0/6 unit 2]
user@host# set vci 6.2
```

   This statement configures VCI 2 on VPI 6.

3. Specify that you want to configure the PPPoE protocol family.

```
[edit interfaces at-1/0/6 unit 2]
user@host# edit family pppoe
```

4. Configure additional PPPoE-specific options for the dynamic subscriber interface.

```
[edit interfaces at-1/0/6 unit 2 family pppoe]
user@host# set duplicate-protection
user@host# set short-cycle-protection lockout-time-min 120 lockout-time-max
   240
user@host# set max-sessions 3
user@host# set max-sessions-vsa-ignore
user@host# set access-concentrator ac-pppoeoa
```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the configuration on logical unit 2 by issuing the **show interfaces at-1/0/6.2** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/6.2
```

```
encapsulation ppp-over-ether-over-atm-llc;
vci 6.2;
family pppoe {
   access-concentrator ac-pppoeoa;
   duplicate-protection;
   max-sessions 3;
   max-sessions-vsa-ignore;
   short-cycle-protection {
      lockout-time-min 120;
      lockout-time-max 240;
   }
}
```

If you are done configuring logical unit 2, enter **commit** from configuration mode.

Configuring the Static PPPoE Subscriber Interface

Step-by-Step Procedure

To configure the static PPPoE subscriber interface:

1. Specify that you want to configure the **pp0** subscriber interface on logical unit 2.

   ```
   [edit]
   user@host# edit interfaces pp0 unit 2
   ```

2. Specify that you want to configure PPP options for the subscriber interface.

   ```
   [edit interfaces pp0 unit 2]
   user@host# edit ppp-options
   ```

3. Configure Challenge Handshake Authentication Protocol (CHAP) authentication for the subscriber interface.

   ```
   [edit interfaces pp0 unit 2 ppp-options]
   user@host# set chap
   user@host# up
   ```

4. Specify that you want to configure PPPoE-specific options.

   ```
   [edit interfaces pp0 unit 2]
   user@host# edit pppoe-options
   ```

5. Associate the PPPoE subscriber interface with the underlying ATM interface.

   ```
   [edit interfaces pp0 unit 2 pppoe-options]
   user@host# set underlying-interface at-1/0/6.2
   ```

6. Configure the router to act as a PPPoE server, also known as a remote access concentrator.

   ```
   [edit interfaces pp0 unit 2 pppoe-options]
   user@host# set server
   user@host# up
   ```

7. Configure the interval for sending keepalive requests.

   ```
   [edit interfaces pp0 unit 2]
   user@host# set keepalives interval 10
   ```

8. Specify that you want to configure the IPv4 (**inet**) protocol family.

   ```
   [edit interfaces pp0 unit 2]
   ```


Copyright © 2013, Juniper Networks, Inc.

```
user@host# edit family inet
```

9. Configure the unnumbered address for the protocol family.

```
[edit interfaces pp0 unit 2 family inet]
user@host# set unnumbered-address lo0.0
```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the static PPPoE subscriber interface configuration by issuing the **show interfaces pp0** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces pp0
unit 2 {
  ppp-options {
    chap;
  }
  pppoe-options {
    underlying-interface at-1/0/6.2;
    server;
  }
  keepalives interval 10;
  family inet {
    unnumbered-address lo0.0;
  }
}
```

If you are done configuring the static PPPoE subscriber interface, enter **commit** from configuration mode.

Verification

To confirm that the static PPPoE subscriber interface pp0.2 is properly configured on ATM underlying interface at-1/0/6.2, perform the following tasks:

- Verifying the ATM Physical Interface Configuration on page 607
- Verifying the Encapsulation, VCI, and PPPoE Options Configuration on Logical Unit 2 on page 608
- Verifying the Static PPPoE Subscriber Interface Configuration on page 609
- Verifying the PPPoE Underlying Interface Configuration on page 609

Verifying the ATM Physical Interface Configuration

Purpose Verify that ATM physical interface at-1/0/6 is properly configured for use with ATM PVCs.

Action From operational mode, issue the **show interfaces at-1/0/6** command.

For brevity, this **show** command output includes only the configuration that is relevant to the at-1/0/6 physical interface. Any other configuration on the system has been replaced with ellipses (...).

```
user@host> show interfaces at-1/0/6
```

```
Physical interface: at-1/0/6, Enabled, Physical link is Down
  Interface index: 179, SNMP ifIndex: 598
  Link-level type: ATM-PVC, MTU: 2048, Clocking: Internal, SDH mode, Speed: OC3,
 Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 00:1f:12:bc:4a:9b
  Last flapped   : 2012-09-19 07:57:59 PDT (07:46:56 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SDH   alarms   : LOL, LOS
  SDH   defects  : LOL, LOS, LOP, BERR-SF, HP-FERF
    VPI 6
      Flags: Active
      Total down time: 0 sec, Last down: Never
  Traffic statistics:
         Input  packets:                    0
         Output packets:                    0
...
```

Meaning **ATM-PVC** in the Link-level Type field indicates that encapsulation for ATM permanent virtual circuits is being used on ATM physical interface at-1/0/6. The **Active** flag for VPI 6 indicates that the virtual path is up and operational.

Verifying the Encapsulation, VCI, and PPPoE Options Configuration on Logical Unit 2

Purpose Verify that the encapsulation, VCI, and PPPoE settings have been properly configured on logical unit 2 (at-1/0/6.2).

Action From operational mode, issue the **show interfaces at-1/0/6.2** command.

```
user@host> show interfaces at-1/0/6.2
  Logical interface at-1/0/6.2 (Index 345) (SNMP ifIndex 1990)
    Flags: Device-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation:
PPPoE-over-ATM-LLC
    Input packets : 0
    Output packets: 0
    Protocol pppoe
      Dynamic Profile: None,
      Service Name Table: None,
      Max Sessions: 3, Max Sessions VSA Ignore: On,
      Duplicate Protection: On, Short Cycle Protection: mac-address,
      AC Name: ac-pppoeoa
    VCI 6.2
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Input packets : 0
        Output packets: 0
```

Meaning **PPPoE-over-ATM-LLC** in the Encapsulation field indicates that logical interface at-1/0/6.2 is properly configured for PPPoE-over-ATM LLC encapsulation. **Protocol pppoe** indicates that the PPPoE protocol family has been properly configured on the logical interface. The **Active** flag for VCI 6.2 indicates that VCI 2 on VPI 6 is up and operational.

 Copyright © 2013, Juniper Networks, Inc.

Verifying the Static PPPoE Subscriber Interface Configuration

Purpose Verify that the static PPPoE subscriber interface (pp0.2) is properly configured.

Action From operational mode, issue the **show interfaces pp0** command.

```
user@host> show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 505
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None

  Logical interface pp0.2 (Index 360) (SNMP ifIndex 1991)
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionDown, Session ID: None,
      Underlying interface: at-1/0/6.2 (Index 345)
    Input packets : 0
    Output packets: 0
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  LCP state: Not-configured
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
  CHAP state: Closed
  PAP state: Closed
    Protocol inet, MTU: 1492
      Flags: Sendbcast-pkt-to-re, Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 100.0.0/24, Local: 100.0.0.1
```

Meaning **PPPoE** in the Link-level type field indicates that PPPoE encapsulation is in use on the **pp0** physical interface. **PPPoE** in the Encapsulation field indicates that PPPoE encapsulation is also in use on the pp0.2 logical subscriber interface. The Underlying interface field indicates that **at-1/0/6.2** is properly configured as the underlying interface for the static PPPoE subscriber interface. **Protocol inet** indicates that the IPv4 protocol family is properly configured on the pp0.2 logical subscriber interface.

Verifying the PPPoE Underlying Interface Configuration

Purpose Verify that the underlying interface is properly configured for static PPPoE-over-ATM subscriber access.

Action From operational mode, issue the **show pppoe underlying-interfaces at-1/0/6.2 extensive** command.

```
user@host> show pppoe underlying-interfaces at-1/0/6.2 extensive
at-1/0/6.2 Index 345
  State: Static, Dynamic Profile: None,
  Max Sessions: 3, Max Sessions VSA Ignore: On,
  Active Sessions: 0,
  Service Name Table: None,
  Duplicate Protection: On, Short Cycle Protection: mac-address,
  AC Name: ac-pppoeoa,
```

```
PacketType                        Sent         Received
  PADI                               0                0
  PADO                               0                0
  PADR                               0                0
  PADS                               0                0
  PADT                               0                0
  Service name error                 0                0
  AC system error                    0                0
  Generic error                      0                0
  Malformed packets                  0                0
  Unknown packets                    0                0
Lockout Time (sec):  Min: 120, Max: 240
  Total clients in lockout: 0
  Total clients in lockout grace period: 0
```

Meaning This command indicates that ATM logical interface at-1/0/6.2 is properly configured as the PPPoE underlying interface. **Static** in the State field indicates that at-1/0/0/.2 is statically configured. The remaining fields display information about the PPPoE-specific interface options configured for the PPPoE underlying interface at the **[edit interfaces at-1/0/6 unit 2 family pppoe]** hierarchy level. The Lockout Time fields, which appear in this command only when you display the **extensive** level of output, show the minimum lockout time (120 seconds) and maximum lockout time (240 seconds) configured for the PPPoE underlying interface.

Related Documentation

- ATM for Subscriber Access Overview on page 581
- Configuring ATM for Subscriber Access on page 587
- Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM on page 593
- Example: Configuring a Static Subscriber Interface for IP Access over ATM on page 610
- Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM on page 616
- Example: Configuring a Static PPP Subscriber Interface over ATM on page 623

Example: Configuring a Static Subscriber Interface for IP Access over ATM

This example illustrates a routed IP-over-ATM (IPoA) configuration that creates a subscriber interface for a static IPv4 interface over a static ATM interface on an MX Series router. The router must have Module Port Concentrator/Modular Interface Card (MPC/MIC) interfaces that use an ATM MIC with small form-factor pluggable transceiver (SFP).

- Requirements on page 611
- Overview on page 611
- Configuration on page 612
- Verification on page 615

 Copyright © 2013, Juniper Networks, Inc.

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.2
- MX Series 3D Universal Edge Router
- ATM MIC with SFP (Model Number MIC-3D-8OC3-2OC12-ATM) and compatible MPC1 or MPC2

Before you begin:

1. Make sure the MX Series router you are using has an ATM MIC with SFP installed and operational.
   - For information about compatible MPCs for the ATM MIC with SFP, see the *MX Series Interface Module Reference*.
   - For information about installing MPCs and MICs in an MX Series router, see the *Hardware Guide* for your MX Series router model.
2. Make sure you understand how to configure and use static ATM interfaces.

   See *ATM Interfaces Overview*.
3. Define the static standard firewall filters (biz-customer-in-filter and biz-customer-out-filter) referenced in the configuration.
   - For information about creating standard firewall filters, see *Guidelines for Configuring Firewall Filters*.
   - For information about applying a firewall filter to an interface, see *Guidelines for Applying Firewall Filters*.

Overview

By using the ATM MIC with SFP and a supported MPC, you can configure the MX Series router to support subscriber access for a statically created IPv4 or IPv6 interface over a static ATM underlying interface. An IPoA configuration enables you to provide access to subscribers on static IPv4 or IPv6 interfaces over an ATM network using ATM Adaptation Layer 5 (AAL5) permanent virtual circuits (PVCs).

NOTE: **IPoA configurations on MX Series routers require static configuration of the IPv4 interface, IPv6 interface, CoS attributes, and firewall fiters. Dynamic configuration is not supported.**

To configure IPoA subscriber access on MX Series routers, you must configure the correct encapsulation type: **atm-snap** for IPoA encapsulation with logical link control (LLC), or **atm-vc-mux** for IPoA encapsulation with virtual circuit (VC) multiplexing. This example configures **atm-vc-mux** as the encapsulation type on the ATM logical interface.

Copyright © 2013, Juniper Networks, Inc.

To provision the ATM AAL5 PVCs for access over the ATM network, you must also configure the virtual path identifiers (VPIs) on the ATM physical interface, and one or more virtual circuit identifiers (VCIs) for each VPI.

In IPoA configurations, the subscriber interfaces correspond to the IPv4 or IPv6 addresses that are on the same network as the statically configured ATM underlying interface. In this IPoA example, the IPv4 address 1.0.0.2 represents the subscriber interface. You can configure the destination address with the **set address 1.0.0.2.254/32 destination 1.0.0.2** statement at the **[edit interfaces at-1/0/3 unit 0 family inet]** hierarchy level.

This example includes the following basic steps to statically configure a single IPv4 subscriber interface over an ATM underlying interface:

1. Configure VPI 0 on ATM physical interface at-1/0/3.
2. Configure ATM VC multiplex encapsulation, VCI 0.39 (VCI 39 on VPI 0), and the following IPv4 (**inet**) protocol family characteristics on logical interface at-1/0/3.0 :
   - IP source address validation (**rpf-check**)
   - Standard input (biz-customer-in-filter) and output (biz-customer-out-filter) firewall filters
   - Interface address 1.0.0.254/32 with destination address 1.0.0.2
3. Configure static access route 200.10.10.0/24 with qualified-next-hop address at-1/0/0.0.

Configuration

To configure a static IPv4 subscriber interface over a static ATM underlying interface, perform these tasks:

- Configuring the ATM Physical Interface on page 613
- Configuring the Static IPv4 Subscriber Interface on Logical Unit 0 on page 613
- Configuring Routing Properties on page 614

CLI Quick Configuration

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

```
# ATM Physical Interface
set interfaces at-1/0/3 atm-options vpi 0
#
# Logical Unit 0
set interfaces at-1/0/3 unit 0 encapsulation atm-vc-mux
set interfaces at-1/0/3 unit 0 vci 0.39
set interfaces at-1/0/3 unit 0 family inet rpf-check
set interfaces at-1/0/3 unit 0 family inet filter input biz-customer-in-filter
set interfaces at-1/0/3 unit 0 family inet filter output biz-customer-out-filter
set interfaces at-1/0/3 unit 0 family inet address 1.0.0.254/32 destination 1.0.0.2
#
# Routing Properties
```


Copyright © 2013, Juniper Networks, Inc.

```
set routing-options access route 200.10.10.0/24 qualified-next-hop at-1/0/0.0
```

Configuring the ATM Physical Interface

Step-by-Step Procedure

To configure the ATM physical interface:

1. Specify that you want to configure ATM-specific options on the physical interface.

```
[edit interfaces at-1/0/3]
user@host# edit atm-options
```

2. Configure one or more VPIs on the physical interface.

```
[edit interfaces at-1/0/3 atm-options]
user@host# set vpi 0
```

Results

From the **[edit]** hierarchy level in configuration mode, confirm the results of the ATM physical interface configuration by issuing the **show interfaces at-1/0/3** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/3
atm-options {
   vpi 0;
}
```

If you are done configuring the ATM physical interface, enter **commit** from configuration mode.

Configuring the Static IPv4 Subscriber Interface on Logical Unit 0

Step-by-Step Procedure

To configure the static IPv4 subscriber interface on logical unit 0:

1. Configure ATM VC multiplex encapsulation on the logical interface.

```
[edit interfaces at-1/0/3 unit 0]
user@host# set encapsulation atm-vc-mux
```

2. Configure the VCI for the logical interface.

```
[edit interfaces at-1/0/3 unit 0]
user@host# set vci 0.39
```

3. Configure the IPv4 (**inet**) protocol family, IPv4 address, and remote (destination) address of the connection.

```
[edit interfaces at-1/0/3 unit 0]
user@host# set family inet address 1.0.0.254/32 destination 1.0.0.2
```

4. Specify that you want to configure additional attributes for the IPv4 protocol family.

```
[edit interfaces at-1/0/3 unit 0]
user@host# edit family inet
```

5. Enable IP source address validation, which checks whether traffic is arriving at the router on an expected path.

```
[edit interfaces at-1/0/3 unit 0 family inet]
```

```
user@host# set rpf-check
```

6. Apply the previously defined standard firewall filters to the logical interface.

```
[edit interfaces at-1/0/3 unit 0 family inet]
user@host# set filter input biz-customer-in-filter
user@host# set filter output biz-customer-out-filter
```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the static subscriber interface configuration on logical unit 0 by issuing the **show interfaces at-1/0/3.0** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/3.0
encapsulation atm-vc-mux;
vci 0.39;
family inet {
   rpf-check;
   filter {
      input biz-customer-in-filter;
      output biz-customer-out-filter;
   }
   address 1.0.0.254/32 {
      destination 1.0.0.2;
   }
}
```

If you are done configuring the static subscriber interface on logical unit 0, enter **commit** from configuration mode.

Configuring Routing Properties

Step-by-Step Procedure To configure static routing properties:

1. Specify that you want to configure protocol-independent routing properties.

```
[edit]
user@host# edit routing-options
```

2. Configure a static access route for routing downstream traffic from the router, and a qualified-next-hop address for routing upstream traffic to the router.

```
[edit routing-options]
user@host# set access route 200.10.10.0/24 qualified-next-hop at-1/0/0.0
```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the static routing properties configuration by issuing the **show routing-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show routing-options
access {
   route 200.10.10.0/24 {
```

 Copyright © 2013, Juniper Networks, Inc.

```
      qualified-next-hop at-1/0/0.0;
    }
  }
```

If you are done configuring the static routing properties, enter **commit** from configuration mode.

Verification

To confirm that the IPoA configuration is working properly, perform the following tasks:

- Verifying the ATM Physical Interface Configuration on page 615
- Verifying the Static Subscriber Interface Configuration on Logical Unit 0 on page 616

Verifying the ATM Physical Interface Configuration

Purpose Verify that the at-1/0/3 physical interface is properly configured for use with ATM PVCs.

Action From operational mode, issue the **show interfaces at-1/0/3** command.

For brevity, this **show** command output includes only the configuration that is relevant to the at-1/0/3 physical interface. Any other configuration on the system has been replaced with ellipses (...).

```
user@host> show interfaces at-1/0/3
Physical interface: at-1/0/3, Enabled, Physical link is Down
  Interface index: 168, SNMP ifIndex: 595
  Link-level type: ATM-PVC, MTU: 2048, Clocking: Internal, SONET mode, Speed:
OC3, Loopback: None,
  Payload scrambler: Enabled
  Device flags   : Present Running Down
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 00:1f:12:bc:4a:98
  Last flapped   : 2012-08-28 07:14:48 PDT (08:28:47 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, LOS
  SONET defects  : LOL, LOS, LOP, BERR-SF, RDI-P
    VPI 0
      Flags: Active
      Total down time: 0 sec, Last down: Never
  Traffic statistics:
        Input  packets:                    0
        Output packets:                    0
  ...
```

Meaning **ATM-PVC** in the Link-level Type field indicates that encapsulation for ATM permanent virtual circuits is being used on ATM physical interface at-1/0/3. The **Active** flag for VPI 0 indicates that the virtual path is up and operational.

Verifying the Static Subscriber Interface Configuration on Logical Unit 0

Purpose Verify that the static subscriber interface on logical unit 0 is properly configured for IPv4 access over ATM.

Action From operational mode, issue the **show interfaces at-1/0/3.0** command.

```
user@host> show interfaces at-1/0/3.0
  Logical interface at-1/0/3.0 (Index 341) (SNMP ifIndex 1984)
    Flags: Device-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: ATM-VCMUX

    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 2040
      Flags: Sendbcast-pkt-to-re, uRPF
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.0.0.2, Local: 1.0.0.254
    VCI 0.39
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Input packets : 0
        Output packets: 0
```

Meaning **ATM-VCMUX** in the Encapsulation field indicates that the logical interface at-1/0/3.0 is properly configured for IPoA encapsulation with VC multiplexing. **Protocol inet** indicates that the IPv4 protocol family has been properly configured on the logical interface. The local address 1.0.0.254 is the IPv4 address of the logical interface. The destination address 1.0.0.2, which is in the same network as the local address, is the IPv4 address of the remote side of the connection and represents the static subscriber interface. The **Active** flag for VCI 0.39 indicates that virtual circuit identifier (VCI) 39 on VPI 0 is up and operational.

Related Documentation

- ATM for Subscriber Access Overview on page 581
- Configuring ATM for Subscriber Access on page 587
- Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM on page 593
- Example: Configuring a Static PPPoE Subscriber Interface over ATM on page 602
- Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM on page 616
- Example: Configuring a Static PPP Subscriber Interface over ATM on page 623

Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM

This example illustrates a bridged IP-over-Ethernet-over-ATM (IPoE-over-ATM) configuration that creates a subscriber interface for IPv4 access over a static ATM interface on an MX Series router. The router must have Module Port Concentrator/Modular

 Copyright © 2013, Juniper Networks, Inc.

Interface Card (MPC/MIC) interfaces that use an ATM MIC with small form-factor pluggable transceiver (SFP).

- Requirements on page 617
- Overview on page 617
- Configuration on page 618
- Verification on page 621

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.2
- MX Series 3D Universal Edge Router
- ATM MIC with SFP (Model Number MIC-3D-8OC3-2OC12-ATM) and compatible MPC1 or MPC2

Before you begin:

1. Make sure the MX Series router you are using has an ATM MIC with SFP installed and operational.
   - For information about compatible MPCs for the ATM MIC with SFP, see the *MX Series Interface Module Reference*.
   - For information about installing MPCs and MICs in an MX Series router, see the *Hardware Guide* for your MX Series router model.
2. Make sure you understand how to configure and use static ATM interfaces.

   See *ATM Interfaces Overview*.
3. Define the static standard firewall filters (biz-customer-in-filter and biz-customer-out-filter) referenced in the configuration.
   - For information about creating standard firewall filters, see *Guidelines for Configuring Firewall Filters*.
   - For information about applying a firewall filter to an interface, see *Guidelines for Applying Firewall Filters*.

Overview

By using the ATM MIC with SFP and a supported MPC, you can configure the MX Series router to support subscriber access for a statically created IPv4 or IPv6 interface over a static ATM underlying interface. An IPoE-over-ATM configuration enables you to provide access to subscribers on static IPv4 or IPv6 interfaces over an underlying ATM interface on an ATM network using ATM Adaptation Layer 5 (AAL5) permanent virtual circuits (PVCs).

Copyright © 2013, Juniper Networks, Inc.

NOTE: **IPoE-over-ATM configurations on MX Series routers require static configuration of the IP interface, ATM interface, CoS attributes, and firewall fiters. Dynamic configuration is not supported.**

To configure bridged IPoE-over-ATM subscriber access on MX Series routers, you must configure Ethernet-over-ATM logical link control (LLC) encapsulation on the ATM underlying interface by including the **encapsulation ether-over-atm-llc** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

To provision the ATM AAL5 PVCs for access over the ATM network, you must also configure the virtual path identifiers (VPIs) on the ATM physical interface, and one or more virtual circuit identifiers (VCIs) for each VPI.

In IPoE-over-ATM configurations, the subscriber interfaces are associated with IPv4 or IPv6 addresses that are mapped to media access control (MAC) addresses. To statically configure Address Resolution Protocol (ARP) table entries that map IP address to MAC addresses, use the **arp** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]** hierarchy level. In this example, the IPv4 address 1.0.50.2, configured with the **set arp 1.0.50.2 mac 00:00:01:02:04:ff publish** statement at the **[edit interfaces at-1/0/2 unit 0 family inet address 1.0.50.254/24]** hierarchy level, represents the subscriber interface.

This example includes the following basic steps to statically configure a single IPv4 subscriber interface over an ATM underlying interface:

1. Configure VPI 0 on ATM physical interface at-1/0/2.
2. Configure Ethernet-over-ATM LLC encapsulation, VCI 0.39 (VCI 39 on VPI 0), and the following IPv4 (**inet**) protocol family characteristics on logical interface at-1/0/2.0 :
   - IPv4 subscriber interface address 1.0.50.254/24
   - Static Address Resolution Protocol (ARP) table entries that provide explicit mappings between IP addresses and MAC addresses
   - IP source address validation (**rpf-check**)
   - Standard input (biz-customer-in-filter) and output (biz-customer-out-filter) firewall filters
3. Configure static access route 200.10.10.0/24 with qualified-next-hop address at-1/0/0.0.

Configuration

To configure a static IPv4 subscriber interface over a static ATM underlying interface, perform these tasks:

- Configuring the ATM Physical Interface on page 619
- Configuring the Static IPv4 Subscriber Interface on Logical Unit 0 on page 619
- Configuring Routing Properties on page 621

 Copyright © 2013, Juniper Networks, Inc.

CLI Quick Configuration

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

```
# ATM Physical Interface
set interfaces at-1/0/2 atm-options vpi 0
#
# Logical Unit 0
set interfaces at-1/0/2 unit 0 encapsulation ether-over-atm-llc
set interfaces at-1/0/2 unit 0 vci 0.39
set interfaces at-1/0/2 unit 0 family inet rpf-check
set interfaces at-1/0/2 unit 0 family inet filter input biz-customer-in-filter
set interfaces at-1/0/2 unit 0 family inet filter output biz-customer-out-filter
set interfaces at-1/0/2 unit 0 family inet address 1.0.50.254/24 arp 1.0.50.2 mac
   00:00:01:02:04:ff
set interfaces at-1/0/2 unit 0 family inet address 1.0.50.254/24 arp 1.0.50.2 publish
#
# Routing Properties
set routing-options access route 200.10.10.0/24 qualified-next-hop at-1/0/0.0
```

Configuring the ATM Physical Interface

Step-by-Step Procedure

To configure the ATM physical interface:

1. Specify that you want to configure ATM-specific options on the physical interface.

   ```
   [edit interfaces at-1/0/2]
   user@host# edit atm-options
   ```

2. Configure one or more VPIs on the physical interface.

   ```
   [edit interfaces at-1/0/2 atm-options]
   user@host# set vpi 0
   ```

Results

From the **[edit]** hierarchy level in configuration mode, confirm the results of the ATM physical interface configuration by issuing the **show interfaces at-1/0/2** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/2
atm-options {
   vpi 0;
}
```

If you are done configuring the ATM physical interface, enter **commit** from configuration mode.

Configuring the Static IPv4 Subscriber Interface on Logical Unit 0

Step-by-Step Procedure

To configure the static IPv4 subscriber interface on logical unit 0:

1. Configure Ethernet-over-ATM LLC encapsulation on the logical interface.

```
[edit interfaces at-1/0/2 unit 0]
user@host# set encapsulation ether-over-atm-llc
```

2. Configure the VCI for the logical interface.

```
[edit interfaces at-1/0/2 unit 0]
user@host# set vci 0.39
```

3. Configure the IPv4 (**inet**) protocol family and address.

```
[edit interfaces at-1/0/2 unit 0]
user@host# set family inet address 1.0.50.254/24
```

4. Specify that you want to configure static ARP table entries to map between IP addresses and MAC addresses.

```
[edit interfaces at-1/0/2 unit 0 family inet]
user@host# edit family inet address 1.0.50.254/24
```

5. Configure IP address 1.0.50.2, which maps to the MAC address, and MAC address 00:00:01:02:04:ff, which maps to the IP address. Include the **publish** option to specify that the router reply to ARP requests for the specified IP address.

```
[edit interfaces at-1/0/2 unit 0 family inet address 1.0.50.254/24]
user@host# set arp 1.0.50.2 mac 00:00:01:02:04:ff publish
user@host# up
```

6. Enable IP source address validation, which checks whether traffic is arriving at the router on an expected path.

```
[edit interfaces at-1/0/2 unit 0 family inet]
user@host# set rpf-check
```

7. Apply the previously defined standard firewall filters to the logical interface.

```
[edit interfaces at-1/0/2 unit 0 family inet]
user@host# set filter input biz-customer-in-filter
user@host# set filter output biz-customer-out-filter
```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the static subscriber interface configuration on logical unit 0 by issuing the **show interfaces at-1/0/2.0** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/2.0
encapsulation ether-over-atm-llc;
vci 0.39;
family inet {
   rpf-check;
   filter {
      input biz-customer-in-filter;
      output biz-customer-out-filter;
   }
   address 1.0.50.254/24 {
      arp 1.0.50.2 mac 00:00:01:02:04:ff publish;
   }
}
```

 Copyright © 2013, Juniper Networks, Inc.

If you are done configuring the static subscriber interface on logical unit 0, enter **commit** from configuration mode.

Configuring Routing Properties

Step-by-Step Procedure

To configure static routing properties:

1. Specify that you want to configure protocol-independent routing properties.

   [edit]
   user@host# **edit routing-options**

2. Configure a static access route for routing downstream traffic from the router, and a qualified-next-hop address for routing upstream traffic to the router.

   [edit routing-options]
   user@host# **set access route 200.10.10.0/24 qualified-next-hop at-1/0/0.0**

Results

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

```
[edit]
user@host# show routing-options
access {
  route 200.10.10.0/24 {
    qualified-next-hop at-1/0/0.0;
  }
}
```

If you are done configuring the static routing properties, enter **commit** from configuration mode.

Verification

To confirm that the IPoE-over-ATM configuration is working properly, perform the following tasks:

- Verifying the ATM Physical Interface Configuration on page 621
- Verifying the Static Subscriber Interface Configuration on Logical Unit 0 on page 622

Verifying the ATM Physical Interface Configuration

Purpose

Verify that the at-1/0/2 physical interface is properly configured for use with ATM PVCs.

Action

From operational mode, issue the **show interfaces at-1/0/2** command.

For brevity, this **show** command output includes only the configuration that is relevant to the at-1/0/2 physical interface. Any other configuration on the system has been replaced with ellipses (...).

```
user@host> show interfaces at-1/0/2
```

```
Physical interface: at-1/0/2, Enabled, Physical link is Down
  Interface index: 175, SNMP ifIndex: 594
  Link-level type: ATM-PVC, MTU: 2048, Clocking: Internal, SDH mode, Speed: OC3,
 Loopback: None,
Payload scrambler: Enabled
  Device flags   : Present Running Down
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 00:1f:12:bc:4a:97
  Last flapped   : 2012-09-06 12:11:39 PDT (05:45:45 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SDH   alarms   : LOL, LOS
  SDH   defects  : LOL, LOS, LOP, BERR-SF, HP-FERF
    VPI 0
      Flags: Active
      Total down time: 0 sec, Last down: Never
  Traffic statistics:
         Input  packets:                    0
         Output packets:                    0
...
```

Meaning **ATM-PVC** in the Link-level Type field indicates that encapsulation for ATM permanent virtual circuits is being used on ATM physical interface at-1/0/2. The **Active** flag for VPI 0 indicates that the virtual path is up and operational.

Verifying the Static Subscriber Interface Configuration on Logical Unit 0

Purpose Verify that the static subscriber interface on logical unit 0 is properly configured for IPoE-over-ATM access.

Action From operational mode, issue the **show interfaces at-1/0/2.0** command.

```
user@host> show interfaces at-1/0/2.0
  Logical interface at-1/0/2.0 (Index 336) (SNMP ifIndex 1983)
    Flags: Device-Down Point-To-Multipoint SNMP-Traps 0x4000 Encapsulation:
Ether-over-ATM-LLC
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 2016
      Flags: Sendbcast-pkt-to-re, uRPF
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.0.50/24, Local: 1.0.50.254, Broadcast: 1.0.50.255
    VCI 0.39
      Flags: Active, Multicast
      Total down time: 0 sec, Last down: Never
        Input packets : 0
        Output packets: 0
```

Meaning **Ether-over-ATM-LLC** in the Encapsulation field indicates that logical interface at-1/0/2.0 is properly configured for Ethernet-over-ATM encapsulation with LLC. **Protocol inet** indicates that the IPv4 protocol family has been properly configured on the logical interface. The destination address 1.0.50/24 identifies the network in which the subscriber interface (1.0.50.2) resides. The **Active** flag for VCI 0.39 indicates that virtual circuit identifier (VCI) 39 on VPI 0 is up and operational.

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- ATM for Subscriber Access Overview on page 581
- Configuring ATM for Subscriber Access on page 587
- Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM on page 593
- Example: Configuring a Static PPPoE Subscriber Interface over ATM on page 602
- Example: Configuring a Static Subscriber Interface for IP Access over ATM on page 610
- Example: Configuring a Static PPP Subscriber Interface over ATM on page 623

Example: Configuring a Static PPP Subscriber Interface over ATM

This example illustrates a PPP-over-ATM (PPPoA) configuration that creates three static PPP logical subscriber interfaces over a static ATM underlying interface on an MX Series router. The router must have Module Port Concentrator/Modular Interface Card (MPC/MIC) interfaces that use an ATM MIC with small form-factor pluggable transceiver (SFP).

- Requirements on page 623
- Overview on page 624
- Configuration on page 625
- Verification on page 629

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.2
- MX Series 3D Universal Edge Router
- ATM MIC with SFP (Model Number MIC-3D-8OC3-2OC12-ATM) and compatible MPC1 or MPC2

Before you begin:

1. Make sure the MX Series router you are using has an ATM MIC with SFP installed and operational.
   - For information about compatible MPCs for the ATM MIC with SFP, see the *MX Series Interface Module Reference*.
   - For information about installing MPCs and MICs in an MX Series router, see the *Hardware Guide* for your MX Series router model.
2. Make sure you understand how to configure and use static ATM interfaces.

   See *ATM Interfaces Overview*.
3. Create the dynamic profile (pppoa-cos-profile) and access profile (pe-B-ppp-clients) referenced in the configuration.

- For information about creating a basic dynamic profile, see "Configuring a Basic Dynamic Profile" on page 673.
- For information about creating a dynamic profile for class of service (CoS) attributes, see "Configuring Traffic Scheduling and Shaping for Subscriber Access" on page 971.
- For information about creating an access profile for PPP Challenge Handshake Authentication Protocol (CHAP) authentication, see *Configuring the PPP Challenge Handshake Authentication Protocol*.

Overview

By using the ATM MIC with SFP and a supported MPC, you can configure an MX Series router to support PPP subscriber access over an ATM network. PPPoA configurations on MX Series routers consist of one or more statically created PPP logical subscriber interfaces over a static ATM underlying interface.

Optionally, you can use dynamic profiles to dynamically or statically apply subscriber services services, such as CoS and firewall filters, to the static PPP logical interface. Configuring CoS and firewall filters in this manner enables you to efficiently and economically provide these services to PPP subscribers accessing the router over an ATM network using ATM Adaptation Layer 5 (AAL5) permanent virtual connections (PVCs). This example uses a previously configured dynamic profile named pppoa-cos-profile to apply traffic scheduling and shaping parameters to logical interface at-1/0/1.2.

To configure PPPoA subscriber access on MX Series routers, you must configure the correct encapsulation type: **atm-ppp-llc** for PPPoA encapsulation with logical link control (LLC), or **atm-ppp-vc-mux** for PPPoA encapsulation with virtual circuit (VC) multiplexing. This example configures **atm-ppp-llc** as the encapsulation type on logical interface at-1/0/1.0, and **atm-ppp-vc-mux** as the encapsulation type on logical interfaces at-1/0/1.1 and at-1/0/1.2.

To provision the ATM AAL5 PVCs for access over the ATM network, you must also configure the virtual path identifiers (VPIs) on the ATM physical interface, and one or more virtual circuit identifiers (VCIs) for each VPI.

In PPPoA configurations, each statically configured logical interface (for example, at-1/0/1.0) corresponds to a PPP logical subscriber interface. This example configures three PPP logical subscriber interfaces over an ATM interface, as follows:

- The ATM physical interface (at-1/0/1) is statically configured with VPI 0 and VPI 2.
- Logical interface at-1/0/1.0 (logical unit 0) is configured with PPP-over AAL5 LLC encapsulation, VCI 0.120 (VCI 120 on VPI 0), PPP-specific options, and the IPv4 protocol family and address.
- Logical interface at-1/0/1.1 (logical unit 1) is configured with PPP-over-AAL5 VC multiplexing encapsulation, VCI 2.120 (VCI 120 on VPI 2), PPP-specific options, and the IPv4 protocol family and address.
- Logical interface at-1/0/1.2 (logical unit 2) is configured with PPP-over-AAL5 VC multiplexing encapsulation, VCI 2.121 (VCI 121 on VPI 2), PPP-specific options, and the

 Copyright © 2013, Juniper Networks, Inc.

IPv4 protocol family and address. The PPP-specific options include applying a dynamic profile named pppoa-cos-profile to the static PPP interface. The pppoa-cos-profile dynamic profile applies traffic scheduling and shaping parameters to the PPP logical subscriber interface.

Configuration

To configure static PPP logical subscriber interfaces over an ATM interface, perform these tasks:

- Configuring the ATM Physical Interface on page 626
- Configuring the Static PPP Subscriber Interface on Logical Unit 0 on page 626
- Configuring the Static PPP Subscriber Interface on Logical Unit 1 on page 627
- Configuring the Static PPP Subscriber Interface on Logical Unit 2 on page 628

CLI Quick Configuration

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

```
# ATM Physical Interface
set interfaces at-1/0/1 atm-options vpi 0
set interfaces at-1/0/1 atm-options vpi 2
#
# Logical Unit 0
set interfaces at-1/0/1 unit 0 encapsulation atm-ppp-llc
set interfaces at-1/0/1 unit 0 vci 0.120
set interfaces at-1/0/1 unit 0 ppp-options chap access-profile pe-B-ppp-clients
set interfaces at-1/0/1 unit 0 ppp-options chap local-name pe-A-at-1/0/1
set interfaces at-1/0/1 unit 0 keepalives interval 5
set interfaces at-1/0/1 unit 0 keepalives up-count 6
set interfaces at-1/0/1 unit 0 keepalives down-count 4
set interfaces at-1/0/1 unit 0 family inet address 192.122.13.13/30
#
# Logical Unit 1
set interfaces at-1/0/1 unit 1 encapsulation atm-ppp-vc-mux
set interfaces at-1/0/1 unit 1 vci 2.120
set interfaces at-1/0/1 unit 1 keepalives interval 6
set interfaces at-1/0/1 unit 1 keepalives up-count 6
set interfaces at-1/0/1 unit 1 keepalives down-count 4
set interfaces at-1/0/1 unit 1 family inet address 192.122.14.13/30
#
# Logical Unit 2
set interfaces at-1/0/1 unit 2 encapsulation atm-ppp-vc-mux
set interfaces at-1/0/1 unit 2 vci 2.121
set interfaces at-1/0/1 unit 2 ppp-options chap access-profile pe-A-ppp-clients
set interfaces at-1/0/1 unit 2 ppp-options chap local-name pe-A-at-1/0/1
set interfaces at-1/0/1 unit 2 ppp-options chap passive
set interfaces at-1/0/1 unit 2 ppp-options dynamic-profile pppoa-cos-profile
set interfaces at-1/0/1 unit 2 keepalives interval 5
set interfaces at-1/0/1 unit 2 keepalives up-count 6
set interfaces at-1/0/1 unit 2 keepalives down-count 4
set interfaces at-1/0/1 unit 2 family inet address 192.122.15.13/30
```

Copyright © 2013, Juniper Networks, Inc.

Configuring the ATM Physical Interface

Step-by-Step Procedure

To configure the ATM physical interface:

1. Specify that you want to configure ATM-specific options on the physical interface.

```
[edit interfaces at-1/0/1]
user@host# edit atm-options
```

2. Configure one or more VPIs on the physical interface.

```
[edit interfaces at-1/0/1 atm-options]
user@host# set vpi 0
user@host# set vpi 2
```

Results

From the **[edit]** hierarchy level in configuration mode, confirm the results of the ATM physical interface configuration by issuing the **show interfaces at-1/0/1** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/1
atm-options {
   vpi 0;
   vpi 2;
}
```

If you are done configuring the ATM physical interface, enter **commit** from configuration mode.

Configuring the Static PPP Subscriber Interface on Logical Unit 0

Step-by-Step Procedure

To configure the static PPP subscriber interface on logical unit 0:

1. Configure PPP-over AAL5 LLC encapsulation on the logical interface.

```
[edit interfaces at-1/0/1 unit 0]
user@host# set encapsulation atm-ppc-llc
```

2. Configure the VCI for the logical interface.

```
[edit interfaces at-1/0/1 unit 0]
user@host# set vci 0.120
```

3. Specify that you want to configure options for PPP CHAP on the logical interface.

```
[edit interfaces at-1/0/1 unit 0]
user@host# edit ppp-options chap
```

4. Assign the previously configured pe-B-ppp-clients access profile to the PPP logical subscriber interface.

```
[edit interfaces at-1/0/1 unit 0 ppp-options chap]
user@host# set access-profile pe-B-ppp-clients
```

5. Configure the local name used by the interface in CHAP challenge and response packets.

Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces at-1/0/1 unit 0 ppp-options chap]
user@host# set local-name "pe-A-at-1/0/1"
user@host# up 2
```

6. Configure the transmission of keepalive messages on the logical interface.

```
[edit interfaces at-1/0/1 unit 0]
user@host# set keepalives interval 5
user@host# set keepalives up-count 6
user@host# set keepalives down-count 4
```

7. Configure the IPv4 (**inet**) protocol family and IP address.

```
[edit interfaces at-1/0/1 unit 0]
user@host# set family inet address 192.122.13.13/30
```

Results

From the **[edit]** hierarchy level in configuration mode, confirm the results of the static PPP subscriber interface configuration on logical unit 0 by issuing the **show interfaces at-1/0/1.0** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/1.0
encapsulation atm-ppp-llc;
vci 0.120;
ppp-options {
   chap {
      access-profile pe-B-ppp-clients;
      local-name pe-A-at-1/0/1;
   }
}
keepalives interval 5 up-count 6 down-count 4;
family inet {
   address 192.122.13.13/30;
}
```

If you are done configuring the PPP logical subscriber interface on logical unit 0, enter **commit** from configuration mode.

Configuring the Static PPP Subscriber Interface on Logical Unit 1

Step-by-Step Procedure

To configure the static PPP subscriber interface on logical unit 1:

1. Configure PPP-over-AAL5 VC multiplexing encapsulation on the logical interface.

```
[edit interfaces at-1/0/1 unit 1]
user@host# set encapsulation atm-ppc-vc-mux
```

2. Configure the VCI for the logical interface.

```
[edit interfaces at-1/0/1 unit 1]
user@host# set vci 2.120
```

3. Configure the transmission of keepalive messages on the logical interface.

```
[edit interfaces at-1/0/1 unit 1]
user@host# set keepalives interval 6
user@host# set keepalives up-count 6
```

```
user@host# set keepalives down-count 4
```

4. Configure the IPv4 (**inet**) protocol family and IP address.

```
[edit interfaces at-1/0/1 unit 1]
user@host# set family inet address 192.122.14.13/30
```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the static PPP subscriber interface configuration on logical unit 1 by issuing the **show interfaces at-1/0/1.1** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/1.1
encapsulation atm-ppp-vc-mux;
vci 2.120;
keepalives interval 6 up-count 6 down-count 4;
family inet {
   address 192.122.14.13/30;
}
```

If you are done configuring the PPP logical subscriber interface on logical unit 1, enter **commit** from configuration mode.

Configuring the Static PPP Subscriber Interface on Logical Unit 2

Step-by-Step Procedure To configure the static PPP subscriber interface on logical unit 2:

1. Configure PPP-over-AAL5 VC multiplex encapsulation on the logical interface.

```
[edit interfaces at-1/0/1 unit 2]
user@host# set encapsulation atm-ppc-vc-mux
```

2. Configure the VCI for the logical interface.

```
[edit interfaces at-1/0/1 unit 2]
user@host# set vci 2.121
```

3. Specify that you want to configure options for PPP CHAP on the logical interface.

```
[edit interfaces at-1/0/1 unit 2]
user@host# edit ppp-options chap
```

4. Assign the previously configured pe-A-ppp-clients access profile to the PPP logical subscriber interface.

```
[edit interfaces at-1/0/1 unit 2 ppp-options chap]
user@host# set access-profile pe-A-ppp-clients
```

5. Configure the local name used by the interface in CHAP challenge and response packets.

```
[edit interfaces at-1/0/1 unit 2 ppp-options chap]
user@host# set local-name “pe-A-at-1/0/1”
```

6. Configure passive mode for CHAP authentication.

```
[edit interfaces at-1/0/1 unit 2 ppp-options chap]
user@host# set passive
```

 Copyright © 2013, Juniper Networks, Inc.

```
user@host# up
```

7. Apply the previously configured pppoa-cos-profile dynamic profile to the PPP logical subscriber interface.

```
[edit interfaces at-1/0/1 unit 2 ppp-options]
user@host# set dynamic-profile pppoa-cos-profile
user@host# up
```

8. Configure the transmission of keepalive messages on the logical interface.

```
[edit interfaces at-1/0/1 unit 2]
user@host# set keepalives interval 5
user@host# set keepalives up-count 6
user@host# set keepalives down-count 4
```

9. Configure the IPv4 (**inet**) protocol family and IP address.

```
[edit interfaces at-1/0/1 unit 2]
user@host# set family inet address 192.122.15.13/30
```

Results From the **[edit]** hierarchy level in configuration mode, confirm the results of the static PPP subscriber interface configuration on logical unit 2 by issuing the **show interfaces at-1/0/1.2** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-1/0/1.2
encapsulation atm-ppp-vc-mux;
vci 2.121;
ppp-options {
   chap {
      access-profile pe-A-ppp-clients;
      local-name pe-A-at-1/0/1;
      passive;
   }
   dynamic-profile pppoa-cos-profile;
}
keepalives interval 5 up-count 6 down-count 4;
family inet {
   address 192.122.15.13/30;
}
```

If you are done configuring the PPP logical subscriber interface on logical unit 2, enter **commit** from configuration mode.

Verification

To confirm that the PPPoA configuration is working properly, perform the following tasks:

- Verifying the ATM Physical Interface Configuration on page 630
- Verifying the Static PPPoA Configuration on Logical Unit 0 on page 630
- Verifying the Static PPPoA Configuration on Logical Unit 1 on page 631
- Verifying the Static PPPoA Configuration on Logical Unit 2 on page 632

Verifying the ATM Physical Interface Configuration

Purpose Verify that the at-1/0/1 physical interface is properly configured for use with ATM PVCs.

Action From operational mode, issue the **show interfaces at-1/0/1** command.

For brevity, this **show** command output includes only the configuration that is relevant to the at-1/0/1 physical interface. Any other configuration on the system has been replaced with ellipses (...).

```
user@host> show interfaces at-1/0/1
Physical interface: at-1/0/1, Enabled, Physical link is Down
  Interface index: 166, SNMP ifIndex: 593
  Link-level type: ATM-PVC, MTU: 2048, Clocking: Internal, SONET mode, Speed:
OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 00:1f:12:bc:4a:96
  Last flapped   : 2012-06-29 15:35:29 PDT (2d 16:24 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, LOS
  SONET defects  : LOL, LOS, LOP, BERR-SF, RDI-P
    VPI 0
      Flags: Active
      Total down time: 0 sec, Last down: Never
  Traffic statistics:
         Input  packets:                    0
         Output packets:                    0
    VPI 2
      Flags: Active
      Total down time: 0 sec, Last down: Never
  Traffic statistics:
         Input  packets:                    0
         Output packets:                    0

  ...
```

Meaning **ATM-PVC** in the Link-level Type field indicates that encapsulation for ATM permanent virtual circuits is being used on ATM physical interface at-1/0/1. The **Active** flag for VPI 0 and VPI 2 indicates that these virtual paths are up and operational.

Verifying the Static PPPoA Configuration on Logical Unit 0

Purpose Verify that the static PPP subscriber interface is properly configured on logical unit 0 (at-1/0/1.0).

Action From operational mode, issue the **show interfaces at-1/0/1.0** command.

```
user@host> show interfaces at-1/0/1.0
Logical interface at-1/0/1.0 (Index 337) (SNMP ifIndex 1979)
    Flags: Device-Down Point-To-Point Inverse-ARP SNMP-Traps 0x4000 Encapsulation:
 ATM-PPP-LLC
    Input packets : 0
```

 Copyright © 2013, Juniper Networks, Inc.

```
    Output packets: 0
  Keepalive settings: Interval 5 seconds, Up-count 6, Down-count 4
  Keepalive: Input: 0 (never), Output: 0 (never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
  CHAP state: Closed
  PAP state: Closed
    Protocol inet, MTU: 2034
      Flags: Sendbcast-pkt-to-re, Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 192.122.13.12/30, Local: 192.122.13.13, Broadcast:
192.122.13.15
    VCI 0.120
      Flags: Active, Inverse-ARP
      Total down time: 0 sec, Last down: Never
      ARP statistics
       Received: 0, Sent: 0, Denied: 0, Operation not supported: 0,
       Bad packet length: 0, Bad protocol: 0, Bad protocol length: 0,
       Bad hardware length: 0, Dropped: 0
       Last received: Never, Last sent: Never
        Input packets : 0
        Output packets: 0
```

Meaning **ATM-PPP-LLC** in the Encapsulation field indicates that logical interface at-1/0/1.0 is properly configured for PPP-over-AAL5 logical link control (LLC) encapsulation. **Protocol inet** indicates that the IPv4 protocol family has been properly configured on the logical interface. The **Active** flag for VCI 0.120 indicates that virtual circuit identifier (VCI) 120 on VPI 0 is up and operational.

Verifying the Static PPPoA Configuration on Logical Unit 1

Purpose Verify that the static PPP subscriber interface is properly configured on logical unit 1 (at-1/0/1.1).

Action From operational mode, issue the **show interfaces at-1/0/1.1** command.

```
user@host> show interfaces at-1/0/1.1

  Logical interface at-1/0/1.1 (Index 338) (SNMP ifIndex 1980)
    Flags: Device-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation:
ATM-PPP-VCMUX
    Input packets : 0
    Output packets: 0
  Keepalive settings: Interval 6 seconds, Up-count 6, Down-count 4
  Keepalive: Input: 0 (never), Output: 0 (never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
  CHAP state: Closed
  PAP state: Closed
    Protocol inet, MTU: 2038
      Flags: Sendbcast-pkt-to-re, Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 192.122.14.12/30, Local: 192.122.14.13, Broadcast:
192.122.14.15
    VCI 2.120
      Flags: Active, Inverse-ARP
```

```
      Total down time: 0 sec, Last down: Never
      ARP statistics
       Received: 0, Sent: 0, Denied: 0, Operation not supported: 0,
       Bad packet length: 0, Bad protocol: 0, Bad protocol length: 0,
       Bad hardware length: 0, Dropped: 0
       Last received: Never, Last sent: Never
        Input packets : 0
        Output packets: 0
```

Meaning **ATM-PPP-VCMUX** in the Encapsulation field indicates that the logical interface at-1/0/1.1 is properly configured for PPP-over-AAL5 VC multiplexing encapsulation. **Protocol inet** indicates that the IPv4 protocol family has been properly configured on the logical interface. The **Active** flag for VCI 2.120 indicates that virtual circuit identifier (VCI) 120 on VPI 2 is up and operational.

Verifying the Static PPPoA Configuration on Logical Unit 2

Purpose Verify that the static PPP subscriber interface is properly configured on logical unit 2 (at-1/0/1.2).

Action From operational mode, issue the **show interfaces at-1/0/1.2** command.

```
user@host> show interfaces at-1/0/1.2
Logical interface at-1/0/1.2 (Index 339) (SNMP ifIndex 1981)
    Flags: Device-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation:
ATM-PPP-VCMUX
    Input packets : 0
    Output packets: 0
  Keepalive settings: Interval 5 seconds, Up-count 6, Down-count 4
  Keepalive: Input: 0 (never), Output: 0 (never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
  CHAP state: Closed
  PAP state: Closed
    Protocol inet, MTU: 2038
      Flags: Sendbcast-pkt-to-re, Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 192.122.15.12/30, Local: 192.122.15.13, Broadcast:
192.122.15.15
    VCI 2.121
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Input packets : 0
        Output packets: 0
```

Meaning **ATM-PPP-VCMUX** in the Encapsulation field indicates that the logical interface at-1/0/1.2 is properly configured for PPP-over-AAL5 VC multiplexing encapsulation. **Protocol inet** indicates that the IPv4 protocol family has been properly configured on the logical interface. The **Active** flag for VCI 2.121 indicates that virtual circuit identifier 121 on VPI 2 is up and operational.

Related Documentation

- ATM for Subscriber Access Overview on page 581
- Configuring ATM for Subscriber Access on page 587

 Copyright © 2013, Juniper Networks, Inc.

- Example: Configuring a Dynamic PPPoE Subscriber Interface over ATM on page 593
- Example: Configuring a Static PPPoE Subscriber Interface over ATM on page 602
- Example: Configuring a Static Subscriber Interface for IP Access over ATM on page 610
- Example: Configuring a Static Subscriber Interface for IP Access over Ethernet over ATM on page 616

Copyright © 2013, Juniper Networks, Inc.

PART 10

Dynamic Profiles for Access and Services

- Dynamic Profiles Overview on page 637
- Configuring Dynamic Profiles on page 673
- Configuring Dynamic Access and Access-Internal Routes for Subscriber Management on page 687
- Dynamic Profile Examples on page 695

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 33

Dynamic Profiles Overview

- Dynamic Profiles Overview on page 638
- Dynamic Variables Overview on page 640
- Junos OS Predefined Variables on page 642
- Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs on page 660
- User-Defined Variables on page 666
- Variable Expressions Overview on page 667
- Access Profiles and Service Profiles Overview on page 670

Copyright © 2013, Juniper Networks, Inc.

Dynamic Profiles Overview

A dynamic profile is a set of characteristics, defined in a type of template, that you can use to provide dynamic subscriber access and services for broadband applications. These services are assigned dynamically to interfaces. The **dynamic-profiles** hierarchy appears at the top level of the CLI hierarchy and contains many Juniper Networks configuration statements that you normally define statically.

Dynamic profile statements appear in the following subhierarchies within the **[edit dynamic-profiles]** hierarchy:

- **class-of-service**
- **firewall**
- **interfaces**
- **predefined-variable-defaults**
- **protocols**
- **routing-instances**
- **routing-options**
- **variables**

This topic covers:

- Dynamic Profile Interface Support on page 638
- What Dynamic Profiles Do on page 638
- How Dynamic Profiles Work on page 639
- Dynamic Profile Version Creation on page 639
- Dynamic Profile Semantic Checking on page 640

Dynamic Profile Interface Support

You can identify subscribers statically or dynamically. To identify subscribers statically, you can reference a static VLAN interface in a dynamic profile. To identify subscribers dynamically, you create variables for demux interfaces that are dynamically created when subscribers log in.

What Dynamic Profiles Do

A dynamic profile acts as a kind of template that enables you to create, update, or remove a configuration that includes client access (for example, interface or protocol) or service (for example, class of service (CoS)) attributes or objects that are created dynamically (for example pseudowire). Using these profiles enables you to consolidate all of the common attributes of a client (and eventually a group of clients) and apply the attributes or dynamically created objects simultaneously.

Copyright © 2013, Juniper Networks, Inc.

How Dynamic Profiles Work

After profiles are created, they reside on the router in a profile library. These profiles can contain various configurations. For example, you can create a client network access configuration, a services activation configuration, or both. When a router interface receives a join message from a client, the router applies the values configured in the specified dynamic profile to that router interface. The profile can contain interface, CoS, and protocol values that are applied directly to the interface. In addition, the dynamic profile can call input or output firewall filters that reside outside of the dynamic profiles hierarchy.

Dynamic Profile Version Creation

You can create new versions of dynamic profiles that are currently in use by subscribers. Dynamic profile version creation is enabled at the **[edit system]** hierarchy level. When enabled, you can create multiple versions of any dynamic profiles on the router. Any subscriber that logs in following a dynamic profile modification uses the latest version of the dynamic profile. Subscribers that are already active continue to use the older version of the dynamic profile until they log out or their session terminates.

When creating versions of dynamic profiles, keep the following in mind:

- You must enable or disable dynamic profile version creation before creating or using any dynamic profiles on the router. Enabling or disabling dynamic profile version creation after dynamic profiles are configured is not supported.

  NOTE: **Before you can enable or disable dynamic profile version creation for a router on which any dynamic profiles are configured, you must first remove all dynamic profiles from the router configuration.**

- Each version of a dynamic profile is stored in the profile database as a new profile.
- The name of the new profile version is derived by appending a string to the original base dynamic profile name. This string contains two dollar sign ($) characters to identify the version field of the profile name. These two characters are followed by numerical characters that represent the “version number” of the dynamic profile (for example, 01).
- The version number of the dynamic profile is automatically generated by the system.
- The dynamic profile that you modify is always stored as the latest version. You cannot create a modified dynamic profile and save it as an earlier version. For example, if you modify version three of a dynamic profile while it is in use, the dynamic profile is saved as version four.
- You can only modify the latest version of a dynamic profile.
- The maximum value for the version number is 99999. However, for each profile, only 10 active versions are supported at a time.
- If the dynamic profile version that you modify is not in use by any subscriber, the profile is overwritten with committed changes without creating a new version.

Copyright © 2013, Juniper Networks, Inc.

- After reaching the 99999th modified version of a dynamic profile, any further modifications to the dynamic profile result in overwriting that final version. If the final version is in use, any modification attempts fail upon commit.
- You can delete a dynamic profile only when none of its versions are in use.
- The dynamic profile version feature supports graceful restart and unified ISSU.

Dynamic Profile Semantic Checking

Variables are applied to dynamic profiles dynamically and cannot be checked with existing CLI commands. Semantic checking validates some variables in dynamic profiles to help identify potential configuration errors.

Semantic checks are performed during commit and during profile instantiation. Commit time checks ensure that variables appear in the correct location within the dynamic profile. Checks performed before profile instantiation ensure that the values that replace the variables are correct. The checks performed on the values include the following:

- Range validation
- Variable type validation
- Existence of variables where they are mandatory
- Variable matching to regular expressions

A commit time check failure results in an error message being displayed and logged in the **/var/log/messages** file and the commit failing. An instantiation failure results in an error being logged in the **/var/log/messages** file and the profile instantiation failing.

Related Documentation

- Configuring a Basic Dynamic Profile on page 673
- Configuring a Dynamic Profile for Client Access on page 679
- Configuring a Dynamic Profile for Various Levels of Services on page 681
- Enabling Dynamic Profiles to use Multiple Versions on page 683
- Dynamic Variables Overview on page 640
- Subscriber Interface Overview on page 757
- *Use Cases for Dynamic Profiles for VPLS Pseudowires*

Dynamic Variables Overview

Variables constitute the dynamic component of a dynamic profile. You use variables in dynamic profiles as placeholders for dynamically obtained or dynamically generated information that the dynamic profiles use to configure subscriber interfaces.

- How Dynamic Variables Work on page 641
- Default Values for Predefined Variables on page 641
- Unique Identifier (UID) for Parameterized Filters on page 641

 Copyright © 2013, Juniper Networks, Inc.

How Dynamic Variables Work

Dynamic variables are data placeholders that you define and place in dynamic profiles. When a particular event occurs on an interface (for example, a DHCP client accesses the interface), the dynamic profiles obtain data to fill these placeholders from one of three possible sources—the interface receiving an incoming client data packet, an externally configured server (for example, RADIUS), or a value associated with each user-configurable variable.

For your convenience, Junos OS provides several predefined variables that you can use within a dynamic profile. Most of these variables relate to interface-specific data obtained directly from the interface that receives an incoming client data packets (for example, interface name, interface unit value, and so on). When a client accesses the interface, the router software extracts the necessary interface data, propagates this data to the dynamic profile, and then uses the dynamic profile to configure the interface for the accessing client.

You define user-defined variables for individual dynamic profiles at the **[dynamic-profiles *profile-name* variables]** hierarchy level. At this hierarchy level, you create an association between a variable value (for example, **$junos-igmp-version**) that appears in the body of the dynamic profile and data associated with that call value that is managed in an externally configured server (for example, a RADIUS VSA managed on a RADIUS server) or defined as a value in the **variables** stanza. When an event occurs on an interface to trigger the instantiation of a dynamic profile for the interface, Junos OS obtains values for each variable from an external server (for example, from RADIUS authentication and authorization VSAs) during the subscriber authentication process. At run time, the variables are replaced by these actual values and are used to configure the subscriber interface.

Default Values for Predefined Variables

You can optionally configure default values for many of the predefined variables. If the external RADIUS server is not available or the VSA does not contain a value for the predefined variable, Junos OS uses the default values.

When a default value is configured for a variable and RADIUS also returns a value, the system uses the value from RADIUS instead.

Unique Identifier (UID) for Parameterized Filters

You can optionally configure a unique identifier (UID) for parameterized filters in dynamic profiles created for services. The generated UIDs enable you to identify and configure separate parameter values for filters with the same variable name. In addition, assigning a UID improves performance of the router.

For service profiles, you can request the generation of an UID for a user-defined variable by including the **uid** statement at the **[dynamic-profiles *profile-name* variables]** hierarchy level. You then reference the variable name in the filter. To enable selection of a particular filter in a dynamic profile that contains multiple variables of the same parameter and criteria type, you must indicate that the variable refers to a UID. To configure, include the **uid-reference** statement at the **[dynamic-profiles *profile-name* variables]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

For example, if the variable **$in-filter** receives the value of “filter1” from RADIUS, the filter definition named $filter is used.

Related Documentation

- Configuring a Basic Dynamic Profile on page 673
- Configuring a Dynamic Profile for Client Access on page 679
- Configuring a Dynamic Profile for Various Levels of Services on page 681
- Junos OS Predefined Variables on page 642
- User-Defined Variables on page 666
- Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs on page 660
- Configuring Predefined Dynamic Variables in Dynamic Profiles on page 675
- Configuring User-Defined Dynamic Variables in Dynamic Profiles on page 676
- Dynamic Profiles Overview on page 638
- Subscriber Interface Overview on page 757
- Example: Firewall Dynamic Profile on page 696
- Example: IGMP Dynamic Profile on page 695
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 90

Junos OS Predefined Variables

Junos OS contains several predefined variables. The dynamic profile obtains and replaces data for these variables from an incoming client data packet and configuration (local and RADIUS). These variables are predefined—you use them in the body of a dynamic profile without first having to define the variables at the **[dynamic-profiles *profile-name* variables]** hierarchy level. Table 59 on page 642 provides a list of predefined variables, their descriptions, and where in the Junos OS hierarchy you can configure them.

Table 59: Junos OS Predefined Variables and Definitions

Variable	Definition
Access and Access-Internal Routes	
$junos-framed-route-cost	Cost metric of an IPv4 access route. You specify this variable at the **[edit dynamic-profiles *profile-name* routing-options access route *address*]** hierarchy level for the **metric** statement.
$junos-framed-route-distance	Distance of an IPv4 access route. You specify this variable at the **[edit dynamic-profiles *profile-name* routing-options access route *address*]** hierarchy level for the **preference** statement.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-framed-route-ip-address-prefix	Route prefix of an IPv4 access route. You specify this variable at the **[edit dynamic-profiles *profile-name* routing-options access]** hierarchy level for the **route** statement.
$junos-framed-route-ipv6-address-prefix	Route prefix of an IPv6 access route. You specify this variable at the **[edit dynamic-profiles *profile-name* routing-options access]** hierarchy level for the **route** statement.
$junos-framed-route-ipv6-cost	Cost metric of an IPv6 access route. You specify this variable with the **metric** statement at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib $junos-ipv6-rib access route $junos-framed-route-ipv6-address-prefix]** hierarchy level.
$junos-framed-route-ipv6-distance	Distance of an IPv6 access route. You specify this variable with the **preference** statement at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib $junos-ipv6-rib access route $junos-framed-route-ipv6-address-prefix]** hierarchy level.
$junos-framed-route-ipv6-nexthop	IPv6 next-hop address of an access route. You specify this variable at the **[edit dynamic-profiles *profile-name* routing-options access route *address*]** hierarchy level for the **next-hop** statement.
$junos-framed-route-ipv6-tag	Tag value of an IPv6 access route. You specify this variable with the **tag** statement at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib $junos-ipv6-rib access route $junos-framed-route-ipv6-address-prefix]** hierarchy level.
$junos-framed-route-nexthop	IPv4 next-hop address of an access route. You specify this variable at the **[edit dynamic-profiles *profile-name* routing-options access route *address*]** hierarchy level for the **next-hop** statement.
$junos-framed-route-tag	Tag value of an IPv4 access route. You specify this variable at the **[edit dynamic-profiles *profile-name* routing-options access route *address*]** hierarchy level for the **tag** statement.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-interface-name	Logical interface of an access-internal route. DHCP or PPP supplies this information when the subscriber logs in. You specify this variable at the **[edit dynamic-profiles *profile-name* routing-options access-internal route *address*]** hierarchy level for the **qualified-next-hop** statement. This variable is also used for creating dynamic IP demux interfaces.
$junos-ipv6-rib	Routing table for an IPv6 access route. You specify this variable with the **rib** statement at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options]** hierarchy level. You can use this variable to specify a nondefault routing instance for the route.
$junos-subscriber-ip-address	IP address of a subscriber identified in an access-internal route. You specify this variable at the **[edit dynamic-profiles *profile-name* routing-options access-internal]** hierarchy level for the **route** statement. This variable is also used for creating dynamic IP demux interfaces.
$junos-subscriber-mac-address	MAC address for a subscriber identified in an access-internal route. You specify this variable at the **[edit dynamic-profiles *profile-name* routing-options access-internal route *address* qualified-next hop *underlying-interface*]** hierarchy level for the **mac-address** statement.
Dynamic Protocols	
$junos-igmp-access-group-name	Specifies the access list to use for the source (S) filter.
$junos-igmp-access-source-group-name	Specifies the access list to use for the source-group (S,G) filter.
$junos-igmp-enable	Ensures that IGMP is not disabled on the interface by an AAA-based authentication and management method (for example, RADIUS). You specify this variable at the **[dynamic-profiles *profile-name* protocols igmp]** hierarchy level for the **interface** statement.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-igmp-immediate-leave	Enables IGMP immediate leave on the interface. You specify this variable at the **[dynamic-profiles *profile-name* protocols igmp]** hierarchy level for the **interface** statement.
$junos-igmp-version	IGMP version configured in a client access profile. Junos OS obtains this information from the RADIUS server when a subscriber accesses the router. The version is applied to the accessing subscriber when the profile is instantiated. You specify this variable at the **[dynamic-profiles *profile-name* protocols igmp]** hierarchy level for the **interface** statement.
$junos-interface-name	Name of the dynamic interface to which the subscriber access client connects. Its use is in dynamically enabling IGMP on the subscriber interface. You specify this variable at the **[dynamic-profiles *profile-name* protocols igmp]** hierarchy level for the **interface** statement. The interface name is derived from concatenating the **$junos-interface-ifd-name** and the **$junos-underlying-interface-unit** variables obtained when a subscriber is created dynamically at the **[dynamic-profiles *profile-name* interfaces]** hierarchy level.
$junos-ipv6-ndra-prefix	Prefix value for the router advertisement interface. Junos OS obtains this information from the RADIUS server when a subscriber accesses the router. The prefix value is applied to the accessing subscriber when the profile is instantiated. You specify this variable at the **[dynamic-profiles *profile-name* protocols router-advertisement interface *$junos-interface-name*]** hierarchy level.
$junos-mld-access-group-name	Specifies the access list to use for the group (G) filter.
$junos-mld-access-source-group-name	Specifies the access list to use for the source-group (S,G) filter.
$junos-mld-enable	Ensures that MLD is not disabled on the interface by an AAA-based authentication and management method (for example, RADIUS). You specify this variable at the **[dynamic-profiles *profile-name* protocols mld]** hierarchy level for the **interface** statement.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-mld-immediate-leave	Enables MLD immediate leave on the interface. You specify this variable at the **[dynamic-profiles *profile-name* protocols mld]** hierarchy level for the **interface** statement.
$junos-mld-version	MLD version configured in a client access profile. Junos OS obtains this information from the RADIUS server when a subscriber accesses the router. The version is applied to the accessing subscriber when the profile is instantiated. You specify this variable at the **[dynamic-profiles *profile-name* protocols mld]** hierarchy level for the **interface** statement.
Dynamic CoS — Traffic-Control Profile Parameters	
$junos-cos-adjust-minimum	Minimum adjusted shaping rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **adjust-minimum** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
$junos-cos-byte-adjust	Byte adjustment value configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **bytes** option with the **overhead-accounting** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
$junos-cos-byte-adjust-cell	Overhead bytes when downstream ATM traffic is in cell-mode. NOTE: Do not configure the $junos-cos-byte-adjust-cell variable when the $junos-cos-byte-adjust variable is configured.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-cos-byte-adjust-frame	Overhead bytes when downstream ATM traffic is in frame-mode. NOTE: Do not configure the $junos-cos-byte-adjust-frame variable when the $junos-cos-byte-adjust variable is configured.
$junos-cos-delay-buffer-rate	Delay-buffer rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **delay-buffer-rate** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
$junos-cos-excess-rate	Excess rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **excess-rate** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
$junos-cos-excess-rate-high	Rate configured for excess high-priority traffic in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **excess-rate-high** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-cos-excess-rate-low	Rate configured for excesslow-priority traffic in a traffic-control profile in a dynamic profile for subscriber access. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **excess-rate-low** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
$junos-cos-guaranteed-rate	Guaranteed rate configured in a traffic-control profile in a dynamic profile Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **guaranteed-rate** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
$junos-cos-guaranteed-rate-burst	Burst size for the guaranteed rate that is configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable with the **burst-size** option in the **guaranteed-rate** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-cos-scheduler-map	Scheduler-map name configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **scheduler-map** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level. NOTE: The scheduler map can be defined dynamically (at the **[edit dynamic-profiles *profile-name* class-of-service scheduler-maps]** hierarchy level) or statically (at the **[edit class-of-service scheduler-maps]** hierarchy level).
$junos-cos-shaping-mode	Shaping mode configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **overhead-accounting** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
$junos-cos-shaping-rate	Shaping rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **shaping-rate** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-cos-shaping-rate-burst	Burst size for the shaping rate configured in a traffic-control profile in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable with the **burst-size** option in the **shaping-rate** statement at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
$junos-cos-shaping-rate-excess-high	Shaping rate configured for excess high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-excess-high-burst	Shaping rate burst size configured for excess high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-excess-low	Shaping rate configured for excess low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-excess-low-burst	Shaping rate burst size configured for excess low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-priority-high	Shaping rate configured for high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-priority-high-burst	Shaping rate burst size configured for high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-cos-shaping-rate-priority-low	Shaping rate configured for low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-priority-low-burst	Shaping rate burst size configured for low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-priority-medium	Shaping rate configured for medium-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-priority-medium-burst	Shaping rate burst size configured for medium-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-traffic-control-profile	Traffic-control profile configured in a dynamic profile for subscriber access. The Junos OS obtains the profile information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **traffic-control-profiles** statement at the **[edit dynamic-profiles *profile-name* class-of-service]** hierarchy level.
Dynamic CoS — Scheduler Parameters	
$junos-cos-scheduler	Name of a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable at the **[edit dynamic-profiles *profile-name* class-of-service schedulers]** hierarchy level.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-cos-scheduler-bs	Buffer size as a percentage of total buffer, specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **buffer-size** statement with the **percent** option at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]** hierarchy level.
$junos-cos-scheduler-pri	Packet-scheduling priority value specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **priority** statement at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]** hierarchy level.
$junos-cos-scheduler-dropfile-any	Name of the drop profile for random early detection (RED) for loss-priority level **any** specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **drop-profile** statement at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name* drop-profile-map loss-priority any protocol any]** hierarchy level. NOTE: The drop profile must be configured statically (at the **[edit class-of-service drop-profiles]** hierarchy level).

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-cos-scheduler-dropfile-high	Name of the drop profile for random early detection (RED) for loss-priority level **high** specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **drop-profile** statement at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name* drop-profile-map loss-priority high protocol any]** hierarchy level. NOTE: The drop profile must be configured statically (at the **[edit class-of-service drop-profiles]** hierarchy level).
$junos-cos-scheduler-dropfile-low	Name of the drop profile for random early detection (RED) for loss-priority level **low** specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **drop-profile** statement at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name* drop-profile-map loss-priority low protocol any]** hierarchy level. NOTE: The drop profile must be configured statically (at the **[edit class-of-service drop-profiles]** hierarchy level) for loss-priority **low**.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-cos-scheduler-dropfile-medium-high	Name of the drop profile for random early detection (RED) for loss-priority level **medium-high** specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **drop-profile** statement at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name* drop-profile-map loss-priority medium-high protocol any]** hierarchy level. NOTE: The drop profile must be configured statically (at the **[edit class-of-service drop-profiles]** hierarchy level).
$junos-cos-scheduler-dropfile-medium-low	Name of the drop profile for random early detection (RED) for loss-priority level **medium-low** specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **drop-profile** statement at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name* drop-profile-map loss-priority medium-low protocol any]** hierarchy level. NOTE: The drop profile must be configured statically (at the **[edit class-of-service drop-profiles]** hierarchy level).
$junos-cos-scheduler-excess-priority	Priority value of the excess rate specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **excess-priority** statement at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-cos-scheduler-excess-rate	Value of the excess rate specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **excess-rate** statement at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]** hierarchy level.
$junos-cos-scheduler-shaping-rate	Value of the shaping rate specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **shaping-rate** statement at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]** hierarchy level.
$junos-cos-scheduler-tx	Transmit rate specified for a scheduler configured in a dynamic profile. Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the dynamic profile is attached. You reference this variable in the **transmit-rate** statement at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]** hierarchy level.
Dynamic Connectivity Fault Management Parameters	
$junos-action-profile	Name of the action profile configured in a dynamic profile.
$junos-ccm-interval	Continuity check interval time configured in a dynamic profile.
$junos-loss-threshold	The number of continuity check messages lost before marking the remote MEP as down, configured in a dynamic profile.
$junos-ma-name-format	Name of the maintenance association name format configured in a dynamic profile.
$junos-md-name-format	Name of the maintenance domain format configured in a dynamic profile.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-ma-name	Name of the maintenance association configured in a dynamic profile.
$junos-md-level	Value of 'Level', configured in a dynamic profile.
$junos-md-name	Name of the maintenance domain configured in a dynamic profile.
$junos-mep-id	The 'MEP' value configured in the dynamic profile.
$junos-remote-mep-id	The 'Remote MEP' value configured in the dynamic profile.
Filters — RADIUS-obtained Policies	
$junos-input-filter	Attaches a filter based on RADIUS VSA 26-10 (Ingress-Policy-Name) or RADIUS attribute 11 (Filter-ID) to the interface.
$junos-input-ipv6-filter	Attaches a filter based on RADIUS VSA 26-106 (IPv6-Ingress-Policy-Name) to the interface.
$junos-output-filter	Attaches a filter based on RADIUS VSA 26-11 (Egress-Policy-Name) to the interface.
$junos-output-ipv6-filter	Attaches a filter based on RADIUS VSA 26-107 (IPv6-Egress-Policy-Name) to the interface.
Subscriber Interfaces — Dynamic Demux Interfaces	
$junos-interface-ifd-name	Name of the device to which the subscriber access client connects. All interfaces are created on this device. Its primary use is in creating single or multiple subscribers on a statically created interface. You specify this variable at the **[dynamic-profiles *profile-name* interfaces]** hierarchy level. When creating a logical underlying interface for a dynamic VLAN demux interface, you must also specify this variable at the **[dynamic-profiles *profile-name* interfaces *demux0* unit *$junos-interface-unit* demux-options underlying-interface]** hierarchy level.
$junos-interface-unit	Creates a unit number assigned to the logical interface. The router supplies this information when the subscriber accesses the network. You specify this variable at the **[dynamic-profiles *profile-name* interfaces *interface-name*]** hierarchy level for the **unit** statement.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-ipv6-address	Selects the IPv6 address of the interface the subscriber uses. You specify this variable at the **[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*]**, **[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*]**, **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family*]**, and **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]** hierarchy level for the **address** statement.
$junos-loopback-interface	Selects the loopback interface the subscriber uses. You specify this variable at the **[dynamic profiles *profile-name* interfaces demux0 unit "$junos-interface-unit" family inet]** hierarchy level for the **unnumbered-address** statement.
$junos-preferred-source-address	Selects the preferred source address associated with the loopback address used for the subscriber. You specify this variable at the **[dynamic profiles *profile-name* interfaces demux0 unit "$junos-interface-unit" family inet unnumbered-address "$junos-loopback-interface"]** hierarchy level for the **preferred-source-address** statement.
$junos-subscriber-ip-address	IP address of the subscriber. You specify this variable at the **[dynamic-profiles *profile-name* interfaces *demux0* unit family *inet* demux-source]** hierarchy level. This variable is also used for creating access-internal routes.
$junos-subscriber-ipv6-address	IPv6 address for subscriber. You specify this variable at the **[dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family inet6 demux-source]** hierarchy level.
$junos-subscriber-ipv6-multi-address	Expands the **demux-source** into multiple addresses; for example, the IPv6 prefix and /128 address for the subscriber. You specify this variable at the **[dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family inet6 demux-source]** hierarchy level.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-underlying-interface	Creates a logical underlying interface for a dynamic IP demux interface. The client logs in on this interface. You specify this variable at the **[dynamic profiles *profile-name* interfaces demux0 unit "$junos-interface-unit" demux-options]** hierarchy level for the **underlying-interface** statement. When configured, the underlying interface is used to determine the **$junos-underlying-interface**, **$junos-underlying-interface-unit**, and **$junos-ifd-name** variables. For example, if the receiving logical interface is ge-0/0/0.1, the **$junos-underlying-interface** variable is set to ge-0/0/0 and the **$junos-underlying-interface-unit** variable is set to 1. This variable is also used for creating access-internal routes.
Subscriber Interfaces — Static VLAN Interfaces	
$junos-interface-ifd-name	Name of the device to which the subscriber access client connects. All interfaces are created on this device. Its primary use is in creating single or multiple subscribers on a statically created interface. You specify this variable at the **[dynamic-profiles *profile-name* interfaces]** hierarchy level.
$junos-underlying-interface-unit	Obtains the unit number for the underlying interface. It specifies the use of the underlying interface for the subscriber. You specify this variable at the **[dynamic-profiles *profile-name* interfaces *$junos-interface-ifd-name*]** hierarchy for the **unit** statement.

Copyright © 2013, Juniper Networks, Inc.

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
Subscriber Interfaces — Dynamic PPPoE Interfaces	
$junos-interface-unit	Specifies the logical unit number when the router dynamically creates a PPPoE logical interface. The **$junos-interface-unit** predefined variable is dynamically replaced with the unit number supplied by the network when the PPPoE subscriber logs in. You specify this variable at the **[edit dynamic-profiles *profile-name* interfaces pp0]** hierarchy level for the **unit** statement.
$junos-underlying-interface	Specifies the name of the underlying Ethernet interface on which the router dynamically creates the PPPoE logical interface. The **$junos-underlying-interface** predefined variable is dynamically replaced with the name of the underlying interface supplied by the network when the PPPoE subscriber logs in. You specify this variable at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" pppoe-options]** hierarchy level for the **underlying-interface** statement.
Subscriber Interfaces — Dynamic Interface Sets	
$junos-interface-set-name	Name of an interface set configured in a dynamic profile. To represent the name of a dynamically created agent circuit identifier (ACI) interface set, use the **$junos-interface-set-name** predefined variable in the **interface-set** statement at the **[edit dynamic-profiles *profile-name* interfaces]** hierarchy level.
$junos-svlan-interface-set-name	Locally generated interface set name for use by dual-tagged VLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is ***physical_interface_name - outer_VLAN_tag***.
Wholesale Networking	

Table 59: Junos OS Predefined Variables and Definitions *(continued)*

Variable	Definition
$junos-interface-name	Name of the dynamic interface to which the subscriber access client connects. Its use is in identifying the subscriber interface. You specify this variable at the **[dynamic-profiles *profile-name* routing-instance *$junos-routing-instance*]** hierarchy level for the **interface** statement. The interface name is derived from concatenating the **$junos-interface-ifd-name** and the **$junos-underlying-interface-unit** variables obtained when a subscriber is created dynamically at the **[dynamic-profiles *profile-name* routing-instance *$junos-routing-instance* interface]** hierarchy level.
$junos-routing-instance	Name of the routing instance to which the subscriber is assigned. This variable triggers a return value from the RADIUS server for Virtual-Router (VSA 26–1). You reference this variable in the statement at the **[dynamic-profiles *profile-name*]** hierarchy level for the **routing-instance** statement.

Related Documentation

- Dynamic Variables Overview on page 640
- Configuring Predefined Dynamic Variables in Dynamic Profiles on page 675
- Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs on page 660
- User-Defined Variables on page 666

Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs

Table 60 on page 660 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.

Table 60: 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

Copyright © 2013, Juniper Networks, Inc.

Table 60: 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
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			
Virtual-Router (26–1)	$junos-routing-instance	Routing instance to which subscriber is assigned	No

Table 60: 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
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
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

Copyright © 2013, Juniper Networks, Inc.

Table 60: 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-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 © 2013, Juniper Networks, Inc.

Table 60: 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 © 2013, Juniper Networks, Inc.

Table 60: 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 © 2013, Juniper Networks, Inc.

Table 60: 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

Related Documentation

- Dynamic Variables Overview on page 640
- Configuring Predefined Dynamic Variables in Dynamic Profiles on page 675
- Junos OS Predefined Variables on page 642

User-Defined Variables

In service profiles, the Junos OS enables you to configure custom variables at the **[dynamic-profiles *profile-name* variables]** hierarchy level and associate those variables with supported RADIUS VSAs. The dynamic profile obtains and replaces data for these variables from an external server (for example, from RADIUS authentication and authorization VSAs) during the subscriber authentication process. At run time, the variables are replaced by these actual values (obtained from default information on the router or from the RADIUS server) and are used to configure the subscriber interface.

For a complete list of supported RADIUS VSAs for which you can create variable associations, see "RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework" on page 90.

Copyright © 2013, Juniper Networks, Inc.

You can also configure the user-defined variables with a default value. The default value provides a standalone configuration for the associated statement or a backup for the statement configuration if the RADIUS server is inaccessible or the VSA attribute does not contain a value.

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring User-Defined Dynamic Variables in Dynamic Profiles on page 676
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 90
- Junos OS Predefined Variables on page 642

Variable Expressions Overview

Junos OS enables you to create expressions—groups of arithmetic operators, string operators, and operands—for use as variables within dynamic profiles. You configure variable expressions at the **[dynamic-profiles *profile-name* variables]** hierarchy level. At run time, the variable expressions are calculated and used as variable values to configure dynamic subscriber interfaces.

When configuring expressions in dynamic profiles, you must adhere to the following rules:

- You can configure expressions only within a variable stanza of a dynamic profile.
- Dynamic profiles that contain expressions must be used only for service activation.
- You can assign expressions only to user-defined variables. You cannot assign expressions to internal variables or predefined variables.
- Expression values are given precedence over default values.
- Entire expressions must be contained within quotation marks (“ ”).
- Strings within the expressions must be quoted within single quotation marks (’ ’) and the single quotation marks can contain only strings.
- White space is treated as a delimiter for all operands and operators. Strings containing spaces that you create within expressions are treated as single strings and include any leading or trailing white space. For example:

```
dynamic-profiles {
  service profile {
    variables {
      scheduler-name;
      video-filter equals “ ‘ Filter 1 ’ ” # Everything within the single quotation marks is
        considered a string, including the leading and trailing white space
    }
  }
}
```

- The expression must be either all arithmetic operators or all string operators; mixing arithmetic operators and string operators is not allowed unless properly converted to the correct type.

- Expressions can refer to other system predefined variables or other user-defined variables. However, no circular referencing between variables is allowed. For example, the following reference is incorrect:

```
dynamic-profiles {
  Service_Profile_1 {
    variables {
      scheduler-name;
      transmit-rate2 equals " ( $transmit-rate1 * 2)/3" # refers to transmit-rate1
      transmit-rate1 equals " ( $transmit-rate2 * 2)/3" # refers to transmit-rate2
    }
  }
}
```

- Any mandatory variable that does not contain a "default" value or an "equals" expression must contain a value as a part of service activation. For example, a RADIUS service VSA like "service-video(value1, value2)" that contains two or fewer mandatory variables in the dynamic service profile definition "service-video" succeeds. The service activation fails if at least one mandatory variable does not have any value associated with it, either through default or equals attribute evaluation.

Table 61 on page 668 lists supported operators and functions you can use to create expressions.

NOTE: **Precedence 5 is the highest level.**

Table 61: Operators and Functions

Operation	Operator	Associativity	Precedence	Action
Arithmetic Addition	+	Left	1	Adds the elements to the right and left of the operator together.
Arithmetic Subtraction	-	Left	1	Subtracts the element to the right of the operator from the element to the left of the operator.
Arithmetic Multiplication	*	Left	2	Multiplies the element to the left of the operator by the element to the right of the operator.
Arithmetic Division	/	Left	2	Divides the element to the left of the operator by the element to the right of the operator.
Arithmetic Modulo	%	Left	2	Divides the element to the left of the operator by the element to the right of the operator and returns the integer remainder. If the element to the left of the operator is less than the element to the right of the operator, the result is the element to the left of the operator.
Concatenation	##	Left	3	Creates a new string by joining the string values to the left of the operator and the values to the right of the operator together.

Copyright © 2013, Juniper Networks, Inc.

Table 61: Operators and Functions *(continued)*

Operation	Operator	Associativity	Precedence	Action
Maximum	max(param1,param2)	Left	4	Takes the maximum of the two values passed as parameters.
Minimum	min(param1,param2)	Left	4	Takes the minimum of the two values passed as parameters.
Round	round(param1)	-	4	Rounds the value to the nearest integer.
Truncate	trunc(param1)	-	4	Truncates a non-integer value to the value left of the decimal point.
Convert to String	toStr(param1)	-	4	Converts the variable inside the parentheses to a null terminated string.
Convert to Integer	toInt(param1)	-	4	Converts the parameter to an integer. A single string or variable is allowed as a parameter.
Random	rand()	-	4	Generates a random numerical value.
If Not Zero	ifNotZero(param1, param2)	Left	4	Returns the second parameter if the first parameter is not zero. Returns NULL if first parameter is zero.
Parentheses	( )	-	5	Groups operands and operators to achieve results different from simple precedence; effectively has the highest precedence.

Expressions are evaluated after variables are populated with values. The evaluation is conducted immediately before profile instantiation and includes value checking. If the computed values are not acceptable, or rules governing expression syntax are broken, the expression evaluation fails, profile instantiation does not occur, and messages are logged to describe the errors.

Table 62 on page 669 lists the possible expression error scenarios and the action taken by the router software.

Table 62: Expression Errors and Actions

Error	Occurance	Action	Variable Value
Parsing error	Commit check phase	Commit fails	not applicable
Circular variable dependency error	Commit check phase	Commit fails	not applicable
Variables inside the expressions are not defined	Commit check phase	Commit fails	not applicable
Divide by zero	Profile Instantiation	Profile instantiation fails	Zero (0)

Copyright © 2013, Juniper Networks, Inc.

Table 62: Expression Errors and Actions *(continued)*

Error	Occurance	Action	Variable Value
Adding string to a number	Profile Instantiation	Profile instantiation fails	Zero (0)
Overflow error	Profile Instantiation	Profile instantiation fails	Undefined
Underflow error	Profile Instantiation	Profile instantiation fails	Undefined

You can also configure the user-defined variables with a default value. The default value provides a standalone configuration for the associated statement or a backup for the statement configuration if the RADIUS server is inaccessible or the VSA attribute does not contain a value.

Related Documentation

- Configuring Variable Expressions in Dynamic Profiles on page 678
- Dynamic Profiles Overview on page 638
- Configuring User-Defined Dynamic Variables in Dynamic Profiles on page 676
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 90
- Junos OS Predefined Variables on page 642

Access Profiles and Service Profiles Overview

Dynamic profiles enable you to configure parameters that enable access and services to subscribers.

Access profiles, also known as client profiles, contain the parameters to grant access and provide basic service to a subscriber during initial login. By configuring the access profile with Junos OS predefined variables, you also enable the service to be activated for those subscribers at login. The RADIUS variables in an access profile map to one or more VSAs.

Service profiles contain parameters that activate or deactivate services for a subscriber. You can apply a service profile with an access profile at login, or apply the service profile separately to modify a service. A service profile maps to a Service VSA.

Functionality Supported in Access and Service Profiles

The types of variables, expressions, and default values that you can use depends on the type of dynamic profile.

Table 63 on page 671 lists the types of variables supported by access profiles and service profiles.

 Copyright © 2013, Juniper Networks, Inc.

Table 63: Types of Variables Supported in Dynamic Profiles

Type of Dynamic Profile	Junos OS Predefined Variable (Local)	Junos OS Predefined Variable (RADIUS)	User-Defined Variable
Access Profile	Yes	Yes	Yes
Service Profile	Yes	No	Yes

Table 64 on page 671 lists the default values, expressions, and unique identifiers supported by access profiles and service profiles.

Table 64: Default Values and Expressions Supported in Dynamic Profiles

Type of Dynamic Profile	Default Values	Expressions	Unique Identifiers
Access Profile	Yes (RADIUS predefined variables only)	No	Yes (Schedulers and Scheduler maps only)
Service Profile	Yes (User-defined variables only)	Yes (Service activation only)	Yes (Firewall filters only)

Related Documentation

- Dynamic Profiles Overview on page 638
- Variable Expressions Overview on page 667
- Unique Identifiers for Firewall Variables in Dynamic Profiles on page 1154

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 34

Configuring Dynamic Profiles

- Configuring a Basic Dynamic Profile on page 673
- Configuring Predefined Dynamic Variables in Dynamic Profiles on page 675
- Configuring Default Values for Predefined Variables in a Dynamic Profile on page 675
- Configuring User-Defined Dynamic Variables in Dynamic Profiles on page 676
- Configuring Variable Expressions in Dynamic Profiles on page 678
- Configuring a Dynamic Profile for Client Access on page 679
- Configuring a Dynamic Profile for Various Levels of Services on page 681
- Configuring Unique Identifiers for Parameterized Filters in Dynamic Profiles on page 682
- Enabling Dynamic Profiles to use Multiple Versions on page 683
- Modifying Dynamic Profiles with Versioning Disabled on page 684

Configuring a Basic Dynamic Profile

This topic describes how to create a basic dynamic profile. A basic profile must contain a profile name and have both an interface variable name (such as **$junos-interface-ifd-name**) included at the **[edit dynamic-profiles *profile-name* interfaces** hierarchy level and logical interface variable name (such as **$junos-underlying-interface-unit** or **$junos-interface-unit**) at the **[edit dynamic-profiles *profile-name* interfaces *variable-interface-name* unit]** hierarchy level.

Before you configure dynamic profiles for initial client access:

1. Configure the necessary router interfaces that you want DHCP clients to use when accessing the network.

   See "Subscriber Interface Overview" on page 757 for information about the types of interfaces you can use with dynamic profiles and how to configure them.

2. Configure all RADIUS values that you want the profiles to use when validating DHCP clients for access to the multicast network.

   See "Configuring RADIUS Server Parameters for Subscriber Access" on page 36

To configure a basic dynamic profile:

1. Name the profile.

```
[edit]
user@host# edit dynamic-profiles basic-profile
```

2. Define the **interface-name** statement with the internal **$junos-interface-ifd-name** variable used by the router to match the interface name of the receiving interface.

```
[edit dynamic-profiles basic-profile]
user@host# edit interfaces $junos-interface-ifd-name
```

3. Define the **unit** statement with the internal variable:

   - When referencing an existing interface, specify the **$junos-underlying-interface-unit** variable used by the router to match the unit value of the receiving interface.
   - When creating dynamic interfaces, specify the **$junos-interface-unit** variable used by the router to generate a unit value for the interface.

```
[edit dynamic-profiles basic-profile interfaces "$junos-interface-ifd-name"]
user@host# set unit $junos-underlying-interface-unit
```

   or

```
[edit dynamic-profiles basic-profile interfaces "$junos-interface-ifd-name"]
user@host# set unit $junos-interface-unit
```

4. If you are using interface sets, you must apply the traffic-control profile to the interface set in the static **[edit class-of-service]** hierarchy. The interface set name must be explicitly referenced in the CoS configuration as part of the static configuration outside of the dynamic profile. The CoS configuration is static and the interface set name must be statically referenced.

   This rule applies to all interface sets except ACI sets.

Related Documentation

- CoS for Interface Sets of Subscribers Overview on page 1019
- Configuring a Dynamic Profile for Client Access on page 679
- Configuring a Dynamic Profile for Various Levels of Services on page 681
- Configuring Predefined Dynamic Variables in Dynamic Profiles on page 675
- Configuring Static Subscriber Interfaces in Dynamic Profiles on page 765
- Configuring VLAN Dynamic Profiles on page 714
- Dynamic Profiles Overview on page 638
- Dynamic Variables Overview on page 640
- Junos OS Predefined Variables on page 642
- Example: Firewall Dynamic Profile on page 696
- Example: IGMP Dynamic Profile on page 695

Copyright © 2013, Juniper Networks, Inc.

Configuring Predefined Dynamic Variables in Dynamic Profiles

This topic discusses how to configure predefined variables in a dynamic profile. The dynamic profile obtains and replaces data for these variables from an incoming client data packet. You can specify these variables in the body of a dynamic profile without having to first define the variables at the **[edit dynamic-profiles *profile-name* variables]** hierarchy level.

Before you configure dynamic variables:

1. Create a basic dynamic profile.

   See "Configuring a Basic Dynamic Profile" on page 673.

2. Ensure that the router hardware is configured in the network to accept subscriber access.

To configure predefined variables in a dynamic profile:

1. Access the desired dynamic profile.

   ```
   [edit]
   user@host# edit dynamic-profiles igmpProfile1
   [edit dynamic-profiles igmpProfile1]
   ```

2. Configure the necessary variables.

   ```
   [edit dynamic-profiles igmpProfile1]
   user@host# set protocols igmp interface $junos-interface-name
   ```

   For a complete list of supported predefined variables, see "Junos OS Predefined Variables" on page 642.

Related Documentation

- Configuring a Basic Dynamic Profile on page 673
- Configuring User-Defined Dynamic Variables in Dynamic Profiles on page 676
- Dynamic Profiles Overview on page 638
- Dynamic Variables Overview on page 640
- Junos OS Predefined Variables on page 642
- Example: Firewall Dynamic Profile on page 696
- Example: IGMP Dynamic Profile on page 695

Configuring Default Values for Predefined Variables in a Dynamic Profile

You can configure default values for the predefined variables that are configured in a dynamic profile. These default values are used when RADIUS does not supply a value.

To configure default values for Junos predefined variables:

1. Specify that you want to configure the dynamic profile.

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

2. Configure the default value for a specific option within a predefined variable.

```
[edit dynamic-profiles profile-name]
user@host# set predefined-variable-defaults predefined-variable variable-option
  default-value
```

NOTE: Do not use the "junos-" prefix when specifying the *predefined-variable*.

Related Documentation

- For a list of predefined variables and options for which you can configure default values, see Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs on page 660
- Junos OS Predefined Variables on page 642
- Dynamic Variables Overview on page 640

Configuring User-Defined Dynamic Variables in Dynamic Profiles

This topic discusses how to configure the user-defined dynamic variables in a dynamic profile. You define user-defined variables for individual dynamic profiles at the **[edit dynamic-profiles *profile-name* variables]** hierarchy level. At this hierarchy level, you create an association between a variable call value (for example, **$junos-igmp-version**) that appears in the body of the dynamic profile and data associated with that call value that is managed in an externally configured server (for example, a RADIUS VSA managed on a RADIUS server) or defined as a default value in the **variables** stanza.

Before you configure dynamic variables:

1. Create a basic dynamic profile.

   See "Configuring a Basic Dynamic Profile" on page 673.

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

   See "Specifying the Authentication and Accounting Methods for Subscriber Access" on page 24.

3. Configure all RADIUS values that you want the profiles to use when validating subscribers.

   See "Configuring RADIUS Server Parameters for Subscriber Access" on page 36

To configure variables in a dynamic profile:

1. Access the **variables** stanza in the desired dynamic profile.

```
user@host# edit dynamic-profiles profile1 variables
[edit dynamic-profiles profile1 variables]
```


Copyright © 2013, Juniper Networks, Inc.

2. Specify a name to identify the variable.

   The variable name can be any alphanumeric value. The name is an association to a variable in the dynamic profile configuration. For example, if you specify a variable name of "**igmp-version**" as the variable name, you must specify the call variable "**$igmp-version**" in the dynamic profile configuration for the statement you want the variable to define.

   ```
   [edit dynamic-profiles igmpProfile1 variables]
   user@host# set igmp-version
   ```

3. Configure the variable using one (or both) of the following methods:

   - Specify a RADIUS attribute and RADIUS tag (when required) for the variable.

     ```
     [edit dynamic-profiles igmpProfile1 variables]
     user@host# set igmp-version radius vendor-id 4874 attribute 78
     ```

   - Configure a default value for the variable.

     ```
     [edit dynamic-profiles igmpProfile1 variables]
     user@host# set igmp-version default-value 3
     ```

   NOTE: You can configure variables by using the RADIUS method, the default value method, or both. If you choose to configure both a RADIUS attribute and a default value for the variable, the RADIUS attribute takes precedence over the default value. However, the dynamic profile applies the default value if the router cannot contact the RADIUS server or if the RADIUS server does not contain a value for the assigned attribute.

4. Configure the call variable in the dynamic profile.

   ```
   [edit dynamic-profiles igmpProfile1]
   user@host# set protocols igmp interface demux0 version $igmp-version
   ```

   NOTE: The call variable must match the name of the variable that you configured in the variables stanza.

Related Documentation

- Dynamic Profiles Overview on page 638
- Dynamic Variables Overview on page 640
- Configuring a Basic Dynamic Profile on page 673
- User-Defined Variables on page 666
- Configuring Predefined Dynamic Variables in Dynamic Profiles on page 675
- Example: Configuring Dynamic Hierarchical Scheduling and Queuing for Subscriber Access on page 1033
- Example: Firewall Dynamic Profile on page 696
- Example: IGMP Dynamic Profile on page 695

Copyright © 2013, Juniper Networks, Inc.

Configuring Variable Expressions in Dynamic Profiles

You can create expressions—groups of arithmetic operators, string operators, and operands—for use as variables within dynamic profiles. These expressions are used as variable values to configure dynamic subscriber interfaces.

To configure dynamic profile variable expressions:

1. Access the dynamic profile for which you want to create variable expressions.

   [edit]
   user@host# **edit dynamic-profiles** ***profile-name***

2. Access the **variables** hierarchy for the dynamic profile.

   [edit dynamic-profiles *profile-name*]
   user@host# **edit variables**

3. Define the variable using the expression operators and operands described in "Variable Expressions Overview" on page 667.

   [edit dynamic-profiles *profile-name* variables]
   user@host# **set** ***expression***

Table 65 on page 678 provides several examples of expressions that you can create using the supported operators and functions.

Table 65: Expression Examples

Example	Description
video-filter equals "' Filter1' "	Assigns the string " Filter1" to the dynamic $video-filter variable.
video-filter2 equals "$video-filter ## ' Filter2' "	Converts dynamic variable "$video-filter" to a string and concatenates the new string with the string " Filter2". The result is the string "$video-filter Filter2" assigned to the $video-filter2 variable.
tempvar equals "120"	Converts "120" to an integer and assigns the integer to the $tempvar variable.
transmit-rate2 equals " ($transmit-rate1 * 2)/3 + $tempvar)"	Multiplies the "transmit-rate1" variable by 2 and divides that value by the sum of 3 and the value of "$tempvar". The result is assigned to the $transmit-rate2 variable.
host-ip equals " '10.0.0.2' "	Assigns the string "10.0.02" to the $host-ip variable.
max-val "max($max1,$max2)"	Assigns the greater of value "max1" or "max2" to the $max-val variable.
min-val "$min($var1,30)"	Assign the smaller of value "var1" and "30" to the $min-val variable.
rounded-var equals "round($var1)"	Rounds off the value of the variable "$var1" to the nearest integer and assigns the value to the $rounded-var variable.
trunc-var equals "trunc(1234.5)"	Truncates the value in parentheses to the left side of the decimal and assigns the resulting value to the $trunc-var variable.

Copyright © 2013, Juniper Networks, Inc.

Table 65: Expression Examples *(continued)*

Example	Description
bwg-shaping-rate equals "$ancp-downstream - ($ancp-downstream % 2 * (1 - $sp-qos-cell-mode))"	Evaluates the expression as per the precedence set in the parentheses.
temp-filter1 equals " 'Filter1' ## toStr($filter)"	Converts the "$filter" variable to a string value and concatenates the converted string to the string "Filter1". The resulting combined string is assigned to the $temp-filter1 variable.

Related Documentation

- Variable Expressions Overview on page 667
- Dynamic Profiles Overview on page 638
- Configuring User-Defined Dynamic Variables in Dynamic Profiles on page 676
- Junos OS Predefined Variables on page 642

Configuring a Dynamic Profile for Client Access

This topic describes how to create a basic dynamic profile that enables DHCP clients to dynamically access the multicast network.

Before you configure dynamic profiles for initial client access:

1. Create a basic dynamic profile.

   See "Configuring a Basic Dynamic Profile" on page 673.

2. Configure the necessary router interfaces that you want accessing DHCP clients to use.

   See "Subscriber Interface Overview" on page 757 for information about the types of interfaces you can use with dynamic profiles and how to configure them.

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

   See "Specifying the Authentication and Accounting Methods for Subscriber Access" on page 24.

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

   See "Configuring RADIUS Server Parameters for Subscriber Access" on page 36

To configure an initial client access dynamic profile:

1. Access an IGMP access profile.

```
user@host# edit dynamic-profiles access-profile
[edit dynamic-profiles access-profile]
user@host#
```

2. Define the IGMP interface with the interface variable.

NOTE: The variable value is replaced by the name of the interface over which the router received the DHCP message.

```
[edit dynamic-profiles access-profile]
user@host# set protocols igmp interface $junos-interface-name
```

3. (Optional) Enable or disable accounting on the IGMP interface.

```
[edit dynamic-profiles access-profile protocols igmp interface "$junos-interface-name"]
user@host# set accounting
```

or

```
[edit dynamic-profiles access-profile protocols igmp interface "$junos-interface-name"]
user@host# set no-accounting
```

NOTE: This statement enables you to override the accounting setting at the IGMP protocol level. For example, if IGMP accounting is enabled at the **[edit protocols igmp interface *interface-name*]** hierarchy level, you can use the **no-accounting** statement to disable accounting for any IGMP interfaces that are dynamically created by the dynamic profile. If IGMP accounting is not enabled at the **[edit protocols igmp interface *interface-name*]** hierarchy level, you can use the **accounting** statement to enable accounting for any IGMP interfaces that are dynamically created by the dynamic profile.

4. Set the IGMP interface to remain enabled.

```
[edit dynamic-profiles access-profile protocols igmp interface "$junos-interface-name"]
user@host# set disable:$junos-igmp-enable
```

NOTE: RADIUS is capable of disabling IGMP. By assigning the enable variable to the **disable** statement, you can ensure that IGMP remains enabled.

5. (Optional) Specify a group policy for the IGMP interface.

```
[edit dynamic-profiles access-profile protocols igmp interface "$junos-interface-name"]
user@host# set group-policy report-reject-policy
```

6. (Optional) Enable immediate leave on the IGMP interface.

```
[edit dynamic-profiles access-profile protocols igmp interface "$junos-interface-name"]
user@host# set immediate-leave:$junos-igmp-immediate-leave
```

7. (Optional) Set the IGMP interface to obtain the IGMP version from RADIUS.

```
[edit dynamic-profiles access-profile protocols igmp interface "$junos-interface-name"]
user@host# set version $junos-igmp-version
```

Related Documentation

- Configuring a Basic Dynamic Profile on page 673

Copyright © 2013, Juniper Networks, Inc.

- Dynamic Profiles Overview on page 638

Configuring a Dynamic Profile for Various Levels of Services

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

Before you configure dynamic profiles for client services:

1. Create a basic dynamic profile.

   See "Configuring a Basic Dynamic Profile" on page 673.

2. Configure a dynamic profile that enables DHCP clients access to the network.

   See "Configuring a Dynamic Profile for Client Access" on page 679

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

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

   See "Specifying the Authentication and Accounting Methods for Subscriber Access" on page 24.

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

   See "Configuring RADIUS Server Parameters for Subscriber Access" on page 36

To configure an initial client access dynamic profile:

1. Access the desired service profile.

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

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

   See "Configuring a Dynamic Profile for Client Access" on page 679.

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

   See "Dynamically Attaching Statically Created Filters for Any Interface Type" on page 1170, "Dynamically Attaching Statically Created Filters for a Specific Interface Family Type" on page 1169, or "Dynamically Attaching Filters Using RADIUS Variables" on page 1172.

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

Related Documentation

- Configuring a Basic Dynamic Profile on page 673

• Dynamic Profiles Overview on page 638

Configuring Unique Identifiers for Parameterized Filters in Dynamic Profiles

This topic discusses how to configure unique identifiers (UID) for parameterized filters in a dynamic profile. The dynamic profile obtains and replaces data for these variables from an incoming client data packet. You can specify these variables in the body of a dynamic profile without having to first define the variables at the **[edit dynamic-profiles *profile-name* variables]** hierarchy level.

Before you configure dynamic variables:

1. Create a basic dynamic profile.

   See "Configuring a Basic Dynamic Profile" on page 673.

2. Ensure that the router hardware is configured in the network to accept subscriber access.

To configure unique identifiers for parameterized filters in a dynamic profile:

1. Access the desired dynamic profile.

   ```
   [edit]
   user@host# edit dynamic-profiles Profile1
   [edit dynamic-profiles Profile1]
   ```

2. Configure the necessary variables.

   ```
   [edit dynamic-profiles service-profile]
   user@host# set variable policer1 uid
   ```

   ```
   [edit dynamic-profiles service-profile]
   user@host# set variables in-filter uid-reference
   ```

Example of a dynamic-profile for parameterized filters:

```
dynamic profile {
  service-profile {
    variable {
      in-filter {
        default-value filter1;
        mandatory;
        uid-reference;
      }
      policer1 {
        uid;
      }
      filter1 {
        uid;
      }
      policer2 {
        uid;
      }
      filter2 {
        uid;
```

 Copyright © 2013, Juniper Networks, Inc.

```
    }
}
```

Related Documentation

- Configuring a Basic Dynamic Profile on page 673
- Configuring User-Defined Dynamic Variables in Dynamic Profiles on page 676
- Dynamic Profiles Overview on page 638
- Dynamic Variables Overview on page 640
- Junos OS Predefined Variables on page 642
- Example: Firewall Dynamic Profile on page 696
- Example: IGMP Dynamic Profile on page 695

Enabling Dynamic Profiles to use Multiple Versions

You can create new versions of dynamic profiles that are currently in use by subscribers. Any subscriber that logs in following a dynamic profile modification uses the latest version of the dynamic profile. Subscribers that are already active continue to use the older version of the dynamic profile until they log out or their session terminates.

NOTE: **You must enable or disable dynamic profile version creation before creating or using any dynamic profiles on the router. Enabling or disabling dynamic profile version creation after dynamic profiles are configured is not supported.**

To configure versioning for dynamic profiles:

1. Access the router system hierarchy level.

   [edit]
   user@host# **edit system**

2. Access the global dynamic profile options.

   [edit system]
   user@host# **edit dynamic-profile-options**

3. Enable version creation for dynamic profiles on the router.

   [edit system dynamic-profile-options]
   user@host# **set versioning**

Related Documentation

- For special considerations when configuring dynamic profile version creation, see Dynamic Profiles Overview on page 638.
- Configuring a Basic Dynamic Profile on page 673

Modifying Dynamic Profiles with Versioning Disabled

You use dynamic profiles to configure large groups of subscribers. However, after you have configured and applied dynamic profiles, be cautious when modifying any dynamic profiles that are in use by active subscribers on the router if you have not enabled the router to use dynamic profile versioning. This section provides guidelines and procedures for modifying existing profiles and applying them to subscriber interfaces if dynamic profile versioning is not enabled on the router.

When modifying dynamic profiles, keep the following considerations in mind:

- Do not modify a dynamic profile when dynamic profile versioning is disabled and the dynamic profile is in use by active subscribers.
- Modifying a dynamic profile when dynamic profile versioning is disabled and when the dynamic profile is in use by active subscribers can lead to unpredictable behavior.

When a dynamic profile is modified and committed when dynamic profile versioning is not enabled, the router:

1. Logs a warning that the profiles are being modified and committed.
2. Determines whether the profile is currently being use by any subscriber.
3. If the profile is in use by a subscriber, the commit fails and the router logs errors to report the conflict.

Juniper Neworks recommends that you only modify dynamic profiles when you have enabled dynamic profile versioning on the router. However, to properly modify a dynamic profile when dynamic profile versioning is disabled on the router:

1. Ensure that no subscribers are using the dynamic profile.
2. Create a new dynamic profile with a different name that contains the desired changes:

   Original Profile

   ```
   profile1 {
     interfaces {
       "$junos-interface-ifd-name" {
         unit "$junos-underlying-interface-unit" {
           family inet {
             filter {
               input "$junos-input-filter";
             }
           }
         }
       }
     }
   }
   ```

   Original DHCP Configuration

   ```
   forwarding-options {
     dhcp-relay {
       traceoptions {
   ```

 Copyright © 2013, Juniper Networks, Inc.

```
            flag all;
        }
        .......
        dynamic-profile profile1;
        ........
    }
}
```

New Profile

```
profile2 {
    interfaces {
        "$junos-interface-ifd-name" {
            unit "$junos-underlying-interface-unit" {
                family inet {
                    filter {
                        input "$junos-input-filter";
                        output "$junos-output-filter; /* added output filter variable */
                    }
                }
            }
        }
    }
}
```

Modified DHCP Configuration

```
forwarding-options {
    dhcp-relay {
        traceoptions {
            flag all;
        }
        .......
        dynamic-profile profile2; /* Name changed from profile1 */
        ........
    }
}
```

3. Commit the configuration containing the modified profile.

   The modified profile is used for any new subscribers that access the router.

Related Documentation

- Configuring a Basic Dynamic Profile on page 673
- Dynamic Profiles Overview on page 638
- Enabling Dynamic Profiles to use Multiple Versions on page 683

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 35

Configuring Dynamic Access and Access-Internal Routes for Subscriber Management

- Access and Access-Internal Routes for Subscriber Management on page 687
- Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 688
- Configuring Dynamic Access Routes for Subscriber Management on page 690
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 692
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 692
- Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 693
- Verifying the Configuration of Access and Access-Internal Routes for Subscriber Management on page 694

Access and Access-Internal Routes for Subscriber Management

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

Access routes typically are used to apply the values of the RADIUS Framed-Route attribute [22] for IPv4 routes and the Framed-IPv6-Route attribute [99] for IPv6 routes. A framed route consists of a prefix that represents a public network behind the CPE, a next-hop gateway, and optional route attributes consisting of a combination of metric, preference, and tag. The only mandatory component of the framed route is the prefix. The next-hop gateway can be specified explicitly in the framed route. Alternatively, the absence of the gateway address implies address 0.0.0.0, which the must resolve using the CPE's IP address. In either case, the convention is that the next-hop gateway is the CPE IP address.

Copyright © 2013, Juniper Networks, Inc.

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

Consider the following rules for resolving the next-hop gateway to determine when each stanza is required:

- If the RADIUS framed route always specifies the next-hop gateway, only the **access** stanza is required in the dynamic profile. The **access-internal** stanza is not required.
- If the RADIUS framed route does not specify the next-hop gateway—as is more common—the variable representing the next-hop, $junos-framed-route-nexthop, defaults to 0.0.0.0. This value implies that the CPE IP address is to be used. For this case, the **access-internal** stanza is required to resolve $junos-framed-route-nexthop to the CPE IP address (represented in the **access-internal** stanza by $junos-subscriber-ip-address) and the logical interface (represented as a qualified next-hop by $junos-interface-name).

BEST PRACTICE: **We recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed route support.**

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 690
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 692
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 692
- RADIUS IETF Attributes Supported by the AAA Service Framework on page 91

Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover

For a subscriber network configured with either nonstop active routing (NSR) or graceful restart, you can configure the router to wait 180 seconds (3 minutes) before removing access routes and access-internal routes for DHCP and PPP subscriber management after a graceful Routing Engine switchover (GRES) takes place.

- Benefits of Delaying Removal of Access Routes and Access-Internal Routes on page 689
- Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes on page 689
- Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes on page 689

Copyright © 2013, Juniper Networks, Inc.

Benefits of Delaying Removal of Access Routes and Access-Internal Routes

The 3-minute delay in removing access routes and access-internal routes after a graceful Routing Engine switchover provides sufficient time for the DHCP client process (jdhcpd), PPP client process (jpppd), or routing protocol process (rpd) to reinstall the access routes and access-internal routes before the router removes the stale routes from the forwarding table. As a result, the risk of traffic loss is minimized because the router always has available subscriber routes for DHCP subscribers and PPP subscribers.

Configuring the router to delay removal of access routes and access-internal routes after a graceful Routing Engine switchover has the following benefits:

- Provides sufficient time to reinstall subscriber routes from the previously active Routing Engine
- Prevents loss of subscriber traffic due to unavailable routes

Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes

In subscriber networks with graceful restart and routing protocols such as BGP and OSPF configured, the router purges any remaining stale access routes and access-internal routes as soon as the graceful restart operation completes, which can occur very soon after completion of the graceful Routing Engine switchover.

Configuring the delay in removing access and access-internal routes after a graceful Routing Engine switchover causes the router to retain the stale routes for a full 180 seconds, which provides sufficient time for the jdhcpd or jpppd client process to reinstall all of the subscriber routes.

Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes

In subscriber networks with nonstop active routing and routing protocols such as BGP and OSPF configured, the routing protocol process (rpd) immediately purges the stale access routes and access-internal routes that correspond to subscriber routes. This removal results in a loss of subscriber traffic.

Configuring the delay in removing access and access-internal routes after a graceful Routing Engine switchover causes the router to retain the stale routes for a full 180 seconds, which prevents potential traffic loss due to unavailable routes.

Related Documentation

- Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 693
- Access and Access-Internal Routes for Subscriber Management on page 687
- Configuring Dynamic Access Routes for Subscriber Management on page 690
- For information about configuring static routes, see the *Junos OS Routing Protocols Library for Routing Devices*

Configuring Dynamic Access Routes for Subscriber Management

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

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

To dynamically configure access routes:

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

   For IPv4:

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

   For IPv6:

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

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

   For IPv4:

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

   For IPv6:

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

3. Configure the metric as a variable.

   For IPv4:

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

   For IPv6:

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

4. Configure the preference as a variable (IPv4 only).

   For IPv4:

Copyright © 2013, Juniper Networks, Inc.

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

For IPv6:

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

5. Configure the tag as a variable (IPv4 only).

   IPv4:

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

   IPv6:

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

BEST PRACTICE: We recommend that you always include the **access-internal** stanza in the dynamic-profile when the **access** stanza is present for framed route support.

Related Documentation

- Access and Access-Internal Routes for Subscriber Management on page 687
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 692
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 692
- Verifying the Configuration of Access and Access-Internal Routes for Subscriber Management on page 694
- RADIUS IETF Attributes Supported by the AAA Service Framework on page 91

Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management

You can dynamically configure access-internal routes. Configuring support for access-internal variables is optional, but it ensures that values from the access-internal variables are used if the next-hop value is missing in the relevant RADIUS attribute—Framed-Route [22] for IPv4 and Framed-IPv6-Route [99] for IPv6.

BEST PRACTICE: **We recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed route support.**

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

To dynamically configure access-internal routes:

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

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

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

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

   NOTE: **Prior to Junos OS Release 10.0, the variable used for qualified-next-hop was $junos-underlying-interface. It is now $junos-interface-name.**

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

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

Related Documentation

- Access and Access-Internal Routes for Subscriber Management on page 687
- Configuring Dynamic Access Routes for Subscriber Management on page 690
- Verifying the Configuration of Access and Access-Internal Routes for Subscriber Management on page 694

Configuring Dynamic Access-Internal Routes for PPP Subscriber Management

You can dynamically configure access-internal routes for PPP subscribers. Configuring support for access-internal variables is optional, but it ensures that values from the access-internal variables are used if the next-hop value is missing in the relevant RADIUS attribute—Framed-Route [22] for IPv4 and Framed-IPv6-Route [99] for IPv6.

Copyright © 2013, Juniper Networks, Inc.

BEST PRACTICE: We recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed route support.

For PPP subscriber interfaces, you do not need to specify the MAC address for access-internal routes.

To dynamically configure access-internal routes for PPP:

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

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

2. Specify the IP address as a variable.

   [edit dynamic-profiles *profile-name* routing-options]
   user@host# edit access-internal route $junos-subscriber-ip-address

3. Specify the qualified-next-hop as a variable.

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

Related Documentation

- Access and Access-Internal Routes for Subscriber Management on page 687
- Configuring Dynamic Access Routes for Subscriber Management on page 690
- Verifying the Configuration of Access and Access-Internal Routes for Subscriber Management on page 694

Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover

In subscriber networks configured with either nonstop active routing (NSR) or graceful restart, you can configure the router to delay for 180 seconds (3 minutes) before removing access routes and access-internal routes for DHCP and PPP subscriber management after a graceful Routing Engine switchover takes place.

To configure the router to delay removal (flushing) of access-routes and access-internal routes after a graceful Routing Engine switchover:

1. Specify that you want to configure subscriber management.

   [edit system services]
   user@host# edit subscriber-management

2. Configure the router to wait 180 seconds before removing access-routes and access-internal routes after a graceful Routing Engine switchover.

   [edit system services subscriber-management]
   user@host# set gres-route-flush-delay

Related Documentation

- Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 688
- Access and Access-Internal Routes for Subscriber Management on page 687
- Configuring Dynamic Access Routes for Subscriber Management on page 690
- For information about configuring static routes, see the *Junos OS Routing Protocols Library for Routing Devices*

Verifying the Configuration of Access and Access-Internal Routes for Subscriber Management

Purpose

View configuration information for access routes and access-internal routes on DHCP and PPP subscribers.

Action

- To display extensive information about access routes and access-internal routes:

  user@host>**show route extensive**

- To display the configuration for access routes:

  user@host>**show route protocol access**

- To display the configuration for access-internal routes:

  user@host> **show route protocol access-internal**

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 690
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 692
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 692

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 36

Dynamic Profile Examples

- Example: IGMP Dynamic Profile on page 695
- Example: Firewall Dynamic Profile on page 696
- Example: Minimum MLPPP Dynamic Profile on page 697
- Example: Minimum PPPoE Dynamic Profile on page 697

Example: IGMP Dynamic Profile

In this example, IGMP is configured for subscriber access using Junos OS predefined variables.

The predefined variables equate to RADIUS settings as follows:

Junos OS Predefined Variable	RADIUS VSA Name	RADIUS Attribute Number
$var-igmp-version	IGMP-Version	26–78
$var-igmp-access-grp	IGMP-Access-Name	26–71
$var-igmp-access-src-grp	IGMP-Access-Src-Name	26–72

```
[edit dynamic-profiles profile-name]
interfaces {
   demux0 {
      unit "$junos-interface-unit" {
         demux-options {
            underlying-interface "$junos-underlying-interface";
         }
         family inet {
            demux-source {
               “$junos-subscriber-ip-address”;
            }
            unnumbered-address lo0.0 preferred-source-address 20.21.0.1;
         }
      }
   }
}
protocols {
   igmp {
```

```
        interface "$junos-interface-name" {
          version "$var-igmp-version";
          group-policy [ "$var-igmp-access-grp" "$var-igmp-access-src-grp" ];
        }
      }
    }
```

NOTE: You must also configure any global IGMP parameters.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679

Example: Firewall Dynamic Profile

In this example, dynamic firewall is configured for subscriber access using Junos IPv4 predefined variables.

The predefined variables equate to RADIUS settings as follows:

Junos OS Predefined Variable	RADIUS VSA Name	RADIUS Attribute Number
$junos-input-filter	Ingress-Policy-Name	26–10
$junos-output-filter	Egress-Policy-Name	26–11

```
dynamic-profiles {
  DynamicFilterProfile {
    interfaces {
      “$junos-interface-ifd-name” {
        unit “$junos-underlying-interface-unit” {
          family inet {
            filter {
              input “$junos-input-filter”;
              output “$junos-output-filter”;
            }
          }
        }
      }
    }
  }
}
```

NOTE: You must also configure any global firewall parameters.

Related Documentation

- Dynamic Firewall Filters Overview on page 1132
- *Configuring Dynamic Firewall Filter Services for Use in Dynamic Profiles*

Copyright © 2013, Juniper Networks, Inc.

Example: Minimum MLPPP Dynamic Profile

This example shows the minimum configuration for a dynamic profile that is used for static LSQ MLPPP bundle interfaces.

```
dynamic-profiles {
  mlppp-profile-1 {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit";
      }
    }
  }
}
```

Related Documentation

- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 363
- Attaching Dynamic Profiles to MLPPP Bundles on page 375

Example: Minimum PPPoE Dynamic Profile

This example shows the minimum configuration for a dynamic profile that is used for static PPPoE interfaces. The configuration must include the **interfaces pp0** stanza.

```
dynamic-profiles {
  ppp-profile-1 {
    interfaces {
      pp0 {
        unit "$junos-interface-unit";
      }
    }
  }
}
```

Related Documentation

- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 363
- Configuring Dynamic Authentication for PPP Subscribers on page 365
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 370

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

PART 11

Dynamic VLANs

- Dynamic VLAN Overview on page 701
- Configuring Dynamic VLANs on page 713
- Dynamic VLAN Examples on page 751

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 37

Dynamic VLAN Overview

- Dynamic 802.1Q VLAN Overview on page 701
- Agent Circuit Identifier-Based Dynamic VLANs Overview on page 702
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704
- Agent Circuit Identifier-Based Dynamic VLANs Bandwidth Management Overview on page 706
- Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ACI Interface Sets on page 709
- Ethernet OAM Support for Service VLANs Overview on page 709

Dynamic 802.1Q VLAN Overview

You can identify VLANs statically or dynamically. You can also configure a mix of static and dynamic VLANs on the same underlying interface.

For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces. Many hosts can be connected to the same Gigabit Ethernet switch, but they cannot be in the same routing or bridging domain.

To identify VLANs statically, you can reference a static VLAN interface in a dynamic profile. To identify subscribers dynamically, you use a variable to specify an 802.1Q VLAN that is dynamically created when a subscriber accesses the network.

Static VLAN Configuration

Static VLAN configuration is not described in this guide. For information about how to statically configure VLANs and stacked VLANs, see the *Junos OS Network Interfaces Library for Routing Devices*. For an example of how to configure static VLANs in a subscriber access network, see the *Junos OS Subscriber Management and Services Library*.

Dynamic VLAN Configuration

You can configure the router to dynamically create VLANs when a client accesses an interface and requests a VLAN ID that does not yet exist. When a client accesses a particular interface, the router instantiates a VLAN dynamic profile that you have associated with the interface. Using the settings in the dynamic profile, the router extracts

Copyright © 2013, Juniper Networks, Inc.

information about the client from the incoming packet (for example, the interface and unit values), saves this information in the routing table, and creates a VLAN or stacked VLAN ID for the client from a range of VLAN IDs that you configure for the interface.

NOTE: **Dynamic VLAN configuration supports the creation of IPv4 (inet), DHCPv4, IPv6 (inet6), and DHCPv6 VLANs.**

Dynamically configuring VLANs or stacked VLANs requires the following general steps:

1. Configure a dynamic profile for dynamic VLAN or dynamic stacked VLAN creation.

   See "Configuring VLAN Dynamic Profiles" on page 714.

2. Associate the VLAN or stacked VLAN dynamic profile with the interface.

   See "Configuring VLAN Interfaces to Use Dynamic Profiles" on page 720.

3. Specify the Ethernet packet type that the VLAN dynamic profile accepts.

   See "Configuring Which VLAN Ethernet Packet Types Dynamic Profiles Can Accept" on page 721.

4. Define VLAN ranges for use by the dynamic profile when creating VLAN IDs.

   See "Configuring VLAN Ranges for Use with Dynamic Profiles" on page 724.

Agent Circuit Identifier-Based Dynamic VLANs Overview

You can configure the router to create dynamic virtual LAN (VLAN) subscriber interfaces for Dynamic Host Configuration Protocol (DHCP) and Point-to-Point Protocol over Ethernet (PPPoE) subscribers based on agent circuit identifier (ACI) information. To use ACI-based dynamic VLAN subscriber interfaces, you must configure them on Modular Port Concentrators/Modular Interface Cards (MPCs/MICs) that face the access side of the network in an MX Series router.

This overview covers the following topics:

- VLAN Architectures and Subscriber Identification on page 702
- ACI-Based Dynamic VLANs and Agent Circuit Identifier Interface Sets on page 703

VLAN Architectures and Subscriber Identification

The following VLAN architectures defined in the DSL Forum Technical Report (TR)-101, Migration to Ethernet-Based DSL Aggregation (April 2006), use different methods to uniquely identify subscribers in Ethernet-based subscriber access networks:

- 1:1 access model using customer VLANs

  Configurations that use the 1:1 access model uniquely identify subscribers by means of VLAN encapsulation; that is, by using the VLAN ID and stacked VLAN (S-VLAN) ID. Subscriber packets received from the access node (such as a digital subscriber line access multiplexer, or DSLAM) that are either single-tagged with a VLAN ID or double-tagged with both an S-VLAN ID and a VLAN ID are examples of 1:1 VLAN

 Copyright © 2013, Juniper Networks, Inc.

configurations because they provide a one-to-one correspondence between an individual subscriber and the VLAN encapsulation.

In the 1:1 VLAN architecture, each customer premises equipment (CPE) or subscriber network has its own dedicated Layer 2 path to the router. Each subscriber network is separated by a customer VLAN (C-VLAN) that is dedicated to a particular customer. The services for each customer are transmitted from the router to the access node by means of that customer's C-VLAN.

The ability to uniquely identify subscribers by means of VLAN encapsulation facilitates delivery of services such as authentication, authorization, and accounting (AAA); class of service (CoS); and filters (policers) to subscribers in a 1:1 VLAN configuration.

- N:1 access model using service VLANs

  Configurations that use the N:1 access model do not uniquely identify subscribers by means of VLAN encapsulation. Instead, these configurations identify subscribers by means of the agent circuit identifier (ACI) information present in DHCP and PPPoE control packets. Subscriber packets received from the access node that are either single-tagged with the same VLAN ID for a group of subscribers or untagged are examples of N:1 VLAN configurations because they provide a many-to-one correspondence between individual subscribers and the VLAN encapsulation.

  In the N:1 VLAN architecture, a service such as video, voice, or data is typically routed to a particular VLAN instead of having multiple services share a single VLAN, as is the case with the 1:1 VLAN architecture. Such VLANs, often referred to as service VLANs, enable service providers to route different services to different routers to functionally separate network services and reduce network complexity.

  Because a VLAN in an N:1 configuration corresponds to a service rather than an individual subscriber, the router uses ACI information in DHCP and PPPoE control packets instead of VLAN encapsulation to uniquely identify subscribers and facilitate application of subscriber-based services.

ACI-Based Dynamic VLANs and Agent Circuit Identifier Interface Sets

For single-tagged, double-tagged, or untagged N:1 configurations that do not use VLAN encapsulation to uniquely identify subscribers, you can configure the router to create dynamic VLAN subscriber interfaces for DHCP and PPPoE subscribers based on ACI information. ACI-based dynamic VLANs uniquely identify subscribers on the router and facilitate application of subscriber-based services, such as CoS and interface-shared filters, to all subscribers that originate from a single household and share the same ACI information.

When you configure an ACI-based dynamic VLAN, the router examines the DHCP and PPPoE control packets to extract the ACI information in order to build a unique dynamic VLAN subscriber interface. The agent-circuit-identifier value is a string that uniquely identifies the subscriber's access node and the digital subscriber line (DSL) on the access node. For DHCP traffic, the agent-circuit-identifier string is in the DHCP option 82 field of DHCP messages. For PPPoE traffic, the agent-circuit-identifier string is in the DSL Forum Agent-Circuit-ID VSA [26-1] of PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) control packets.

Copyright © 2013, Juniper Networks, Inc.

Configuring ACI-based dynamic VLAN subscriber interfaces is particularly useful in configurations with multiple DHCP and PPPoE subscriber sessions per household. Because DHCP and PPPoE control traffic sent to the router from the same household has the same unique agent-circuit-identifier string, the router groups these DHCP and PPPoE subscriber interfaces in the same ACI interface set. An *ACI interface set* is a logical collection of subscriber interfaces that originate at the same household or on the same access-loop port. Grouping subscriber interfaces into ACI interface sets enables unique subscriber identification and facilitates application of subscriber-based services, such as class of service (CoS) and interface-shared filters, on a per-household basis.

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration on page 747
- Clearing Agent Circuit Identifier Interface Sets on page 748

Agent Circuit Identifier-Based Dynamic VLANs Components Overview

You can configure ACI-based dynamic VLAN subscriber interfaces on Modular Port Concentrators/Modular Interface Cards (MPCs/MICs) that face the access side of the network in an MX Series router.

This overview describes the components of an ACI-based dynamic VLAN configuration, from top to bottom of the interface stack:

- ACI-Based Dynamic Subscriber Interface on page 704
- Dynamic ACI Interface Set on page 704
- Dynamic or Static Underlying VLAN Interface on page 705
- Static Physical Interface on page 705

ACI-Based Dynamic Subscriber Interface

You must create a dynamic profile to define either a dynamic PPPoE subscriber interface for PPPoE subscriber sessions, or a dynamic IP demultiplexer (IP demux) subscriber interface for DHCP subscriber sessions. The router automatically creates (instantiates) the subscriber interface when a DHCP or PPPoE subscriber logs in on the associated underlying VLAN interface associated with the dynamic profile that defines the ACI interface set.

Dynamic ACI Interface Set

The dynamic ACI interface set, which is the primary component of an ACI-based dynamic VLAN configuration, groups the DHCP and PPPoE subscriber sessions that belong to a particular household and share a common unique agent-circuit-identifier value. The router creates one ACI interface set per household to facilitate application of subscriber-based services, such as CoS and interface-shared filters, to all subscribers in the household.

 Copyright © 2013, Juniper Networks, Inc.

You must create a dynamic profile to define the ACI interface set, which is represented in the profile by the Junos OS predefined dynamic variable **$junos-interface-set-name**. When a DHCP or PPPoE subscriber accesses the router on a particular interface, the router obtains the agent-circuit-identifier information from the DHCP or PPPoE control packets transmitted on that interface and dynamically creates the ACI interface set when the first subscriber from that household logs in.

Dynamic or Static Underlying VLAN Interface

After you define the ACI interface set, you must configure the underlying VLAN interface to enable creation of dynamic VLAN subscriber interfaces based on ACI information. You can configure the underlying VLAN interface either dynamically (with a dynamic profile) or statically.

ACI-based dynamic VLAN configurations support the following underlying VLAN interface types:

- Gigabit Ethernet
- 10-Gigabit Ethernet
- VLAN demux (demux0)

NOTE: When you configure an underlying VLAN interface to support creation of ACI-based dynamic VLANs, we recommend that you use this underlying interface only for subscriber interfaces that contain agent-circuit-identifier information in their DHCP or PPPoE control packets. If the router receives DHCP or PPPoE control packets without agent-circuit-identifier information on an underlying VLAN interface configured for ACI-based dynamic VLANs, the associated subscriber interfaces might not instantiate successfully.

Static Physical Interface

ACI-based dynamic VLAN configurations support the following physical interface types:

- Gigabit Ethernet
- 10-Gigabit Ethernet
- Aggregated Ethernet

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Overview on page 702
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration on page 747
- Clearing Agent Circuit Identifier Interface Sets on page 748

Copyright © 2013, Juniper Networks, Inc.

Agent Circuit Identifier-Based Dynamic VLANs Bandwidth Management Overview

A router in a subscriber access network ensures class of service (CoS) for dynamic subscriber interfaces. An MX Series router with Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces ensures that subscribers receive an adequate minimum bandwidth, referred to as the *guaranteed rate*, and maximum bandwidth, referred to as the *shaping rate*. For dynamic VLAN subscriber interfaces based on agent circuit identifier (ACI) information, you can shape the bandwidth either at a per-household level for a dynamic ACI interface set, or at a per-subscriber level for a dynamic VLAN subscriber interface associated with an ACI interface set.

To help you manage bandwidth more efficiently and economically for ACI-based dynamic VLAN subscriber interfaces for PPPoE subscribers, you can configure the router to use specific PPPoE vendor-specific attributes (VSAs) found in PPPoE control packets to adjust the CoS shaping-rate and overhead-accounting attributes for dynamic ACI interface sets and their associated ACI-based dynamic VLAN subscriber interfaces.

This overview covers the following topics:

- CoS Shaping Rate Adjustment on page 706
- CoS Overhead Accounting Adjustment on page 707
- Dynamic Profiles and Adjustment of CoS Shaping Rate and Overhead Accounting on page 707
- Guidelines for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting on page 708

CoS Shaping Rate Adjustment

The CoS shaping rate adjustment is based on the value of the Actual-Data-Rate-Downstream DSL Forum VSA [26-130] found in PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) control packets for PPPoE traffic. The Actual-Data-Rate-Downstream VSA contains the actual downstream data rate, in bits per second, of the subscriber's synchronized digital subscriber line (DSL) link.

To configure the router to use the Actual-Data-Rate-Downstream VSA to adjust the CoS shaping-rate attribute, include the **vendor-specific-tags** statement with the **actual-data-rate-downstream** option at the **[edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]** hierarchy level in either the dynamic profile that defines the ACI interface set or the dynamic profile that configures the associated dynamic PPPoE (**pp0**) subscriber interface.

When you enable this feature, the value of the Actual-Data-Rate-Downstream VSA overrides the **shaping-rate** value configured at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles]** hierarchy level only if the Actual-Data-Rate-Downstream VSA value is less than the **shaping-rate** value configured with the CLI.

Copyright © 2013, Juniper Networks, Inc.

CoS Overhead Accounting Adjustment

The CoS overhead accounting adjustment is based on the value of the Access-Loop-Encapsulation DSL Forum VSA [26-144] found in PADI and PADR control packets for PPPoE traffic. The Access-Loop-Encapsulation VSA identifies the encapsulation used by the subscriber associated with the digital subscriber line access multiplexer (DSLAM) access loop from which requests are initiated.

The value of the Data Link subfield in the Access-Loop-Encapsulation VSA determines the overhead accounting mode in use on the access loop. If the Data Link subfield value is 0 (ATM Adaptation Layer 5, or AAL5), the access loop uses cell-mode encapsulation. If the Data Link subfield value is 1 (Ethernet), the access loop uses frame-mode encapsulation.

In subscriber access networks where the router passes downstream ATM traffic to Ethernet interfaces, the different Layer 2 encapsulations between the router and the PPPoE Intermediate Agent on the DSLAM make managing the bandwidth of downstream ATM traffic difficult. Using the Access-Loop-Encapsulation VSA to shape traffic based on frames or cells enables the router to adjust the overhead-accounting attribute in order to apply the correct downstream rate for the subscriber.

To configure the router to use the Access-Loop-Encapsulation VSA to adjust the CoS overhead-accounting attribute, include the **vendor-specific-tags** statement with the **access-loop-encapsulation** option at the **[edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]** hierarchy level in either the dynamic profile that defines the ACI interface set or the dynamic profile that configures the associated dynamic PPPoE (**pp0**) subscriber interface.

When you enable this feature, the value of the Access-Loop-Encapsulation VSA always overrides the **overhead-accounting** value configured at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles]** hierarchy level.

Dynamic Profiles and Adjustment of CoS Shaping Rate and Overhead Accounting

When you configure the router to use one or both of the Actual-Data-Rate-Downstream VSA value and Access-Loop-Encapsulation VSA value to adjust the CoS shaping rate and overhead accounting attributes, respectively, the router adjusts these attributes when the dynamic ACI interface set is created and the router receives the PADI and PADR packets from the first subscriber interface belonging to the ACI interface set.

You can configure CoS adjustment based on either or both VSAs in either or both of the following dynamic profiles:

- To configure adjustment of the CoS shaping rate and overhead accounting on a per-household basis, use the dynamic profile that defines the dynamic ACI interface set.
- To configure adjustment of the CoS shaping rate and overhead accounting on a per-subscriber basis, use the dynamic profile that defines the ACI-based dynamic PPPoE (**pp0**) subscriber interface associated with the ACI interface set.

Copyright © 2013, Juniper Networks, Inc.

Table 66 on page 708 summarizes how the dynamic profile in which you configure CoS adjustment for ACI-based dynamic VLANs using one or both VSAs affects the router behavior.

Table 66: CoS Adjustment in Dynamic Profiles for ACI Interface Sets and ACI-Based Subscriber Interfaces

VSAs Specified in ACI Interface Set Dynamic Profile	VSAs Specified in PPPoE Subscriber Interface Dynamic Profile	Result
Yes	No	Router adjusts specified CoS attributes only for dynamic ACI interface set
No	Yes	Router adjusts specified CoS attributes only for ACI-based dynamic PPPoE subscriber interface
Yes	Yes	Router adjusts specified CoS attributes for both dynamic ACI interface set and ACI-based dynamic PPPoE subscriber interface
No	No	Router does not adjust CoS attributes for either the dynamic ACI interface set or the ACI-based dynamic PPPoE subscriber interface

Guidelines for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting

You can also configure the router to use the Actual-Data-Rate-Downstream VSA and Access-Loop-Encapsulation VSA values in PPPoE control packets to adjust the CoS shaping rate and overhead accounting attributes, respectively, for dynamic subscriber interfaces *not* associated with dynamic ACI interface sets.

With the exception of the constraints described in "Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ACI Interface Sets" on page 709, most of the guidelines and restrictions that apply to this feature for use with non–ACI-based dynamic subscriber interfaces also apply to its use for dynamic ACI interface sets and their associated ACI-based dynamic VLAN subscriber interfaces.

Related Documentation

- Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 1074
- Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs on page 742
- Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ACI Interface Sets on page 709

 Copyright © 2013, Juniper Networks, Inc.

Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ACI Interface Sets

The following restrictions apply when you configure the router to use the Actual-Data-Rate-Downstream VSA and Access-Loop-Encapsulation vendor-specific attribute (VSA) values in PPPoE control packets to adjust the CoS shaping rate and overhead accounting attributes, respectively, for dynamic ACI interface sets and their associated agent circuit identifier (ACI)-based dynamic VLAN subscriber interfaces:

- You cannot configure adjustment of CoS shaping rate and overhead accounting attributes based on Actual-Data-Rate-Downstream VSA and Access-Loop-Encapsulation VSA values that the router receives from the following sources:
  - RADIUS servers
  - Access Node Control Protocol (ANCP) access loop information
  - Dynamic Host Configuration Protocol (DHCP) discovery packets
- You cannot use this feature to report information about the PPPoE VSA values to RADIUS.
- You cannot use this feature to configure CoS adjustment of upstream data traffic on a dynamic ACI interface set.

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Bandwidth Management Overview on page 706
- Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 1074
- Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs on page 742

Ethernet OAM Support for Service VLANs Overview

You can enable propagation of the Ethernet IEEE 802.1ag Operation, Administration, and Maintenance (OAM) state of a static single-tagged service VLAN (S-VLAN) to a dynamic or static double-tagged customer VLAN (C-VLAN) and, by extension, to the subscriber interfaces configured on the C-VLAN. The static S-VLAN logical interface must be configured on a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet physical interface.

Propagation of the S-VLAN OAM state to associated C-VLANs ensures that when the OAM state of the S-VLAN link is down, the associated C-VLANs and all subscriber interfaces configured on the C-VLANs are brought down as well.

- Ethernet OAM Support for Service VLANs Terms and Acronyms on page 710
- Components of Ethernet OAM Support for Service VLANs on page 710

Copyright © 2013, Juniper Networks, Inc.

- How Ethernet OAM Support for Service VLANs Works on page 711
- Restrictions for Using Ethernet OAM Support for Service VLANs on page 712

Ethernet OAM Support for Service VLANs Terms and Acronyms

Table 67 on page 710 defines the basic terms and acronyms used in this discussion of Ethernet OAM support for service VLANs.

Table 67: Ethernet OAM Support for Service VLANs Terms and Acronyms

Term	Definition
CFM	Connectivity fault management. Provides end-to-end monitoring of an Ethernet network that can be made up of one or more service instances. Junos OS supports Ethernet IEEE 802.1ag CFM.
Continuity check protocol	A feature of Ethernet IEEE 802.1ag CFM that provides fault detection within a maintenance association.
C-VLAN	Customer VLAN. A dynamic or static double-tagged logical interface that has both an outer VLAN tag (corresponding to the S-VLAN) and an inner VLAN tag (corresponding to the C-VLAN). In a 1:1 subscriber network access model, dedicated C-VLANs provide a one-to-one correspondence between an individual subscriber and the VLAN encapsulation.
OAM	Operation, Administration, and Maintenance. A set of Ethernet connectivity specifications and functions providing connectivity monitoring, fault detection and notification, fault verification, fault isolation, loopback, and remote defect identification. Ethernet interfaces on MX Series routers support the IEEE 802.1ag standard for OAM.
S-VLAN	Service VLAN. A static single-tagged logical interface that has only one outer VLAN tag (corresponding to the S-VLAN). In an N:1 subscriber network access model, S-VLANs are dedicated to a particular service, such as video, voice, or data, instead of to a particular subscriber. Because an S-VLAN is typically shared by many subscribers within the same household or in different households, it provides a many-to-one correspondence between individual subscribers and the VLAN encapsulation.
VLAN	Virtual local area network. A logical group of network devices that appear to be on the same local area network, regardless of their physical location.

Components of Ethernet OAM Support for Service VLANs

Ethernet OAM support for S-VLANs involves the following components:

- Physical interface—On MX Series routers with Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces, you can enable propagation of the S-VLAN OAM state to a C-VLAN on Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet physical interfaces.
- S-VLAN—To enable propagation of the S-VLAN Ethernet OAM state to associated C-VLANs and subscriber interfaces, you must configure the static single-tagged S-VLAN logical interface to run the Ethernet IEEE 802.1ag CFM continuity check protocol.

 Copyright © 2013, Juniper Networks, Inc.

- C-VLAN—The C-VLAN is a dynamic or static double-tagged logical interface that has the same S-VLAN (outer) tag as the static single-tagged S-VLAN logical interface. If propagation of the S-VLAN OAM state to the C-VLAN is enabled on the physical interface, the router brings down the C-VLAN and its associated subscriber interfaces when the CFM continuity check protocol detects that the OAM state of the underlying S-VLAN is down.
- Subscriber interfaces—Propagation of the S-VLAN Ethernet OAM state to associated C-VLANs and subscriber interfaces applies to all dynamic or static DHCP, IP demultiplexing (IP demux), and PPPoE subscriber interfaces configured on the C-VLAN.

How Ethernet OAM Support for Service VLANs Works

To enable propagation of the Ethernet OAM state of the S-VLAN to associated C-VLANs and subscriber interfaces, use the **oam-on-svlan** statement when you configure a Gigabit Ethernet (ge), 10-Gigabit Ethernet (xe), or aggregated Ethernet (ae) physical interface.

If Ethernet IEEE 802.1ag CFM is properly configured on the S-VLAN logical interface, including the **oam-on-svlan** statement for these Ethernet interfaces causes the router to bring down both of the following when the CFM continuity check protocol detects that the OAM state of the S-VLAN logical interface is down:

- All dynamic or static double-tagged C-VLAN logical interfaces that have the same S-VLAN (outer) tag as the S-VLAN logical interface on which they are configured.
- All dynamic or static DHCP, IP demux, and PPPoE logical subscriber interfaces configured on the associated C-VLANs.

To illustrate how Ethernet OAM support for S-VLANs works, consider the following sample configuration on a Gigabit Ethernet physical interface:

- Gigabit Ethernet physical interface ge-1/0/3 configured with the **svlan-on-oam** statement.
- Static single-tagged S-VLAN logical interface ge-1/0/3.0, which has a single S-VLAN outer tag, VLAN ID 600.
- Ethernet OAM CFM protocol configured on the static S-VLAN logical interface. The CFM configuration includes an action profile with the **interface-down** default action to bring down the C-VLAN and dynamic subscriber interfaces when the continuity check protocol detects that the Ethernet OAM state of S-VLAN interface ge-1/0/3.0 is down.
- Static double-tagged C-VLAN logical interface ge-1/0/3.100, which has an S-VLAN outer tag, VLAN ID 600, and a C-VLAN inner tag, VLAN ID 1.
- Static PPPoE subscriber interfaces configured on C-VLAN interface ge-1/0/3.100.

Because the S-VLAN and C-VLAN logical interfaces in this example have the same S-VLAN outer tag (VLAN ID 600), the router brings down the C-VLAN interface and the PPPoE logical subscriber interfaces when the CFM continuity check detects that the OAM status of S-VLAN interface ge-1/0/3.0 is down.

Copyright © 2013, Juniper Networks, Inc.

Restrictions for Using Ethernet OAM Support for Service VLANs

Ethernet OAM support for S-VLANs is *not currently supported* for use with any of the following:

- Dynamically configured S-VLAN logical interfaces
- S-VLAN trunk interfaces
- C-VLAN trunk interfaces

Related Documentation

- Configuring Ethernet OAM Support for Service VLANs with Double-Tagged Customer VLANs on page 743
- *IEEE 802.1ag OAM Connectivity Fault Management Overview*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 38

Configuring Dynamic VLANs

- Configuring VLAN Dynamic Profiles on page 714
- Configuring VLAN Interfaces to Use Dynamic Profiles on page 720
- Configuring Which VLAN Ethernet Packet Types Dynamic Profiles Can Accept on page 721
- Configuring an Authentication Password for VLAN or Stacked VLAN Ranges on page 723
- Configuring VLAN Ranges for Use with Dynamic Profiles on page 724
- Configuring Dynamic Authentication for VLAN Interfaces on page 728
- Configuring VLAN Interface Username Information for AAA Authentication on page 729
- Option 82 Suboptions in Authentication Usernames for Autosense VLANs on page 730
- Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs on page 731
- Automatically Removing VLANs with No Subscribers on page 731
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Defining Agent Circuit Identifier Interface Sets on page 735
- Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information on page 737
- Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information on page 739
- Configuring Dynamic VLAN Subscriber Interfaces Based on Agent Circuit Identifier Information on page 740
- Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs on page 742
- Configuring Ethernet OAM Support for Service VLANs with Double-Tagged Customer VLANs on page 743
- Verifying and Managing Dynamic VLAN Configuration on page 747
- Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration on page 747
- Clearing Agent Circuit Identifier Interface Sets on page 748

Copyright © 2013, Juniper Networks, Inc.

Configuring VLAN Dynamic Profiles

Creating dynamic single-tag VLANs or stacked (dual-tag) VLANs requires the use of dynamic profiles. The dynamic profile automatically references the VLAN interface and creates the interface unit and the necessary VLAN IDs for each new single-tag VLAN or stacked VLAN.

NOTE: VLAN dynamic profiles do not support user-defined variables. Use only Junos VLAN predefined variables when configuring VLAN dynamic profiles. See "Dynamic Variables Overview" on page 640 for information about dynamic variables.

- Configuring a VLAN Dynamic Profile for Creating Single-Tag VLANs Using Standard TPID Values on page 714
- Configuring a VLAN Dynamic Profile for Creating Single-Tag VLANs Using Any TPID Values on page 716
- Configuring a Stacked VLAN Dynamic Profile on page 717
- Configuring a VLAN Dynamic Profile That Associates VLAN Interfaces with Separate Routing Instances on page 718

Configuring a VLAN Dynamic Profile for Creating Single-Tag VLANs Using Standard TPID Values

You can configure a VLAN dynamic profile to create single-tag VLANs that accept only standard TPID values (a TPID value of 0x8100) by using the **vlan-id** statement and the **$junos-vlan-id** variable.

NOTE: This procedure configures a dynamic profile that accepts only TPID values of 0x8100. To configure a VLAN dynamic profile for creating VLANs using any TPID values, see "Configuring a VLAN Dynamic Profile for Creating Single-Tag VLANs Using Any TPID Values" on page 716.

Before you begin:

- Configure the dynamic profile.

  See "Configuring a Basic Dynamic Profile" on page 673.

To configure a dynamic VLAN profile:

1. Ensure that the VLAN dynamic profile uses the **$junos-interface-ifd-name** variable for the dynamic interface and the **$junos-interface-unit** variable for the interface unit.
2. (Optional) To support dynamic demux interfaces, enable them using the demux-source statement.
   a. For IPv4 demux interfaces, specify **inet** as the source type.

Copyright © 2013, Juniper Networks, Inc.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set demux-source inet
```

b. For IPv6 demux interfaces, specify **inet6** as the source type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set demux-source inet6
```

3. (Optional) To configure the router to respond to any ARP request, specify the proxy-arp statement.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set proxy-arp
```

4. Specify that you want to use dynamic VLAN IDs in the dynamic profile.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set vlan-id $junos-vlan-id
```

When the dynamic profile is instantiated, the variable is dynamically replaced with a VLAN ID within the VLAN range specified at the **[interfaces]** hierarchy level.

5. Define the unit family type.

a. For IPv4 interfaces, specify the **inet** family type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set family inet
```

b. For IPv6 interfaces, specify the **inet6** family type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set family inet6
```

6. (Optional) Enable IP and MAC address validation for dynamic demux interfaces in a dynamic profile.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit" family inet]
user@host# set mac-validate loose
```

7. Specify the unnumbered address and preferred source address.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit" family inet]
user@host# set unnumbered-address lo.0 preferred-source-address 192.0.16.1
```

Configuring a VLAN Dynamic Profile for Creating Single-Tag VLANs Using Any TPID Values

You can configure a VLAN dynamic profile to create single-tag VLANs that accept any TPID values by configuring the **vlan-tags** statement and the **$junos-vlan-id** variable.

NOTE: For procedures to configure a VLAN dynamic profile for creating single-tag VLANs that use only standard TPID values (a TPID value of 0x8100), see "Configuring a VLAN Dynamic Profile for Creating Single-Tag VLANs Using Standard TPID Values" on page 714.

Before you begin:

- Configure the dynamic profile.

  See "Configuring a Basic Dynamic Profile" on page 673.

To configure a dynamic VLAN profile:

1. Ensure that the VLAN dynamic profile uses the **$junos-interface-ifd-name** variable for the dynamic interface and the **$junos-interface-unit** variable for the interface unit.
2. (Optional) To support dynamic demux interfaces, enable them using the demux-source statement.
   a. For IPv4 demux interfaces, specify **inet** as the source type.

      ```
      [edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
        "$junos-interface-unit"]
      user@host# set demux-source inet
      ```

   b. For IPv6 demux interfaces, specify **inet6** as the source type.

      ```
      [edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
        "$junos-interface-unit"]
      user@host# set demux-source inet6
      ```

3. (Optional) To configure the router to respond to any ARP request, specify the proxy-arp statement.

   ```
   [edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
     "$junos-interface-unit"]
   user@host# set proxy-arp
   ```

4. Specify that you want to use dynamic VLAN IDs in the dynamic profile.

   ```
   [edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
     "$junos-interface-unit"]
   user@host# set vlan-tags outer $junos-vlan-id
   ```

   The variable is dynamically replaced with both the TPID value and a VLAN ID within the VLAN range specified at the **[interfaces]** hierarchy level.

5. Define the unit family type.
   a. For IPv4 interfaces, specify the **inet** family type.

Copyright © 2013, Juniper Networks, Inc.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set family inet
```

b. For IPv6 interfaces, specify the **inet6** family type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set family inet6
```

6. (Optional) Enable IP and MAC address validation for dynamic demux interfaces in a dynamic profile.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit" family inet]
user@host# set mac-validate loose
```

7. Specify the unnumbered address and preferred source address.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit" family inet]
user@host# set unnumbered-address lo.0 preferred-source-address 192.0.16.1
```

Configuring a Stacked VLAN Dynamic Profile

You can configure a dynamic profile for creating stacked 802.1Q VLANs.

Before you begin:

- Configure the dynamic profile.

  See "Configuring a Basic Dynamic Profile" on page 673.

To configure a stacked VLAN dynamic profile:

1. Ensure that the VLAN dynamic profile uses the **$junos-interface-ifd-name** variable for the dynamic interface and the **$junos-interface-unit** variable for the interface unit.
2. (Optional) To support dynamic demux interfaces, enable them using the demux-source statement.

   a. For IPv4 demux interfaces, specify **inet** as the source type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set demux-source inet
```

   b. For IPv6 demux interfaces, specify **inet6** as the source type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set demux-source inet6
```

3. (Optional) To configure the router to respond to any ARP request, specify the proxy-arp statement.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set proxy-arp
```

4. Specify the outer VLAN ID variable.

```
[edit dynamic-profiles STACKED-VLAN-PROF1 interfaces "$junos-interface-ifd-name"
  unit "$junos-interface-unit"]
user@host# set vlan-tags outer $junos-stacked-vlan-id
```

   The variable is dynamically replaced with an outer VLAN ID within the VLAN range specified at the **[interfaces]** hierarchy level.

5. Specify the inner VLAN ID variable.

```
[edit dynamic-profiles STACKED-VLAN-PROF1 interfaces "$junos-interface-ifd-name"
  unit "$junos-interface-unit"]
user@host# set vlan-tags inner $junos-vlan-id
```

   The variable is dynamically replaced with an inner VLAN ID within the VLAN range specified at the **[interfaces]** hierarchy level.

6. Define the unit family type.

   a. For IPv4 interfaces, specify the **inet** family type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set family inet
```

   b. For IPv6 interfaces, specify the **inet6** family type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set family inet6
```

7. (Optional) Enable IP and MAC address validation for dynamic demux interfaces in a dynamic profile.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit" family inet]
user@host# set mac-validate loose
```

8. Specify the unnumbered address and preferred source address.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit" family inet]
user@host# set unnumbered-address lo.0 preferred-source-address 192.0.16.1
```

Configuring a VLAN Dynamic Profile That Associates VLAN Interfaces with Separate Routing Instances

You can configure a VLAN dynamic profile that dynamically creates underlying VLAN interfaces and associates these interfaces with dynamically created routing instances. The VLAN interface is created in the default logical system (LS) for a specific routing instance as defined by VSA 26–1 (Virtual-Router) on the AAA server (for example, RADIUS server).

To configure a dynamic VLAN profile using routing instances:

1. Name the profile.

```
[edit]
```

 Copyright © 2013, Juniper Networks, Inc.

```
user@host# edit dynamic-profiles VLAN_PROFILE_RI
```

2. Specify that you want to dynamically create routing instances on the default logical system.

```
[edit dynamic-profiles VLAN_PROFILE_RI]
user@host# edit routing-instances $junos-routing-instance
```

3. Define the routing instance **interface** statement with the internal **$junos-interface-name** variable used by the router to match the interface name of the receiving interface.

```
[edit dynamic-profiles VLAN_PROFILE_RI routing-instances “$junos-routing-instance”]
user@host# set interface $junos-interface-name
```

4. Define the dynamic profile **interfaces** statement with the internal **$junos-interface-ifd-name** variable.

```
[edit dynamic-profiles VLAN_PROFILE_RI]
user@host# edit interfaces $junos-interface-ifd-name
```

5. Define the **unit** statement with the internal **$junos-interface-unit** variable used by the router to generate a unit value for the interface.

```
[edit dynamic-profiles VLAN_PROFILE_RI interfaces "$junos-interface-ifd-name"]
user@host# edit unit $junos-interface-unit
```

6. To support dynamic demux interfaces, enable them using the demux-source statement.

   a. For IPv4 demux interfaces, specify **inet** as the source type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set demux-source inet
```

   b. For IPv6 demux interfaces, specify **inet6** as the source type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set demux-source inet6
```

7. (Optional) To configure the router to respond to any ARP request, specify the proxy-arp statement.

```
[edit dynamic-profiles VLAN_PROFILE_RI interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set proxy-arp
```

8. Specify that you want to use dynamic VLAN IDs in the dynamic profile.

```
[edit dynamic-profiles VLAN_PROFILE_RI interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set vlan-id $junos-vlan-id
```

   The variable is dynamically replaced with both the TPID value and a VLAN ID within the VLAN range specified at the **[interfaces]** hierarchy level.

9. Define the unit family type.

   a. For IPv4 interfaces, specify the **inet** family type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
```

```
user@host# set family inet
```

b. For IPv6 interfaces, specify the **inet6** family type.

```
[edit dynamic-profiles VLAN-PROF1 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set family inet6
```

10. (Optional) Enable IP and MAC address validation for dynamic demux interfaces in a dynamic profile.

```
[edit dynamic-profiles VLAN_PROFILE_RI interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit" family inet]
user@host# set mac-validate loose
```

11. Specify the unnumbered address to dynamically apply a loopback interface.

NOTE: You can optionally specify the preferred source address. This option is included in the step.

```
[edit dynamic-profiles VLAN_PROFILE_RI interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit" family inet]
user@host# set unnumbered-address $junos-loopback-interface
  preferred-source-address 192.0.16.1
```

Configuring VLAN Interfaces to Use Dynamic Profiles

You can configure an interface to use a single-tag VLAN or stacked (dual-tag) VLAN dynamic profile when creating dynamic VLANs. The dynamic profile assigns a VLAN ID to each VLAN dynamically created over the interface by using the single-tag VLAN and stacked VLAN ranges configured for the VLAN interface. You can configure VLAN interfaces to use dynamic profiles in the following ways:

- Associating a Single-Tag VLAN Dynamic Profile with an Interface on page 720
- Associating a Stacked VLAN Dynamic Profile with an Interface on page 721

Associating a Single-Tag VLAN Dynamic Profile with an Interface

Before you begin:

- Configure the VLAN dynamic profile.

  See "Configuring a Basic Dynamic Profile" on page 673.

To associate a single-tag VLAN dynamic profile with an interface:

1. Access the interface that you want to use for creating VLANs.

```
[edit]
user@host# edit interfaces ge-1/0/0
```

2. Edit the **auto-configure** stanza to automatically configure VLANs.

```
[edit interfaces ge-1/0/0]
user@host# edit auto-configure
```

 Copyright © 2013, Juniper Networks, Inc.

3. Edit the ***vlan-ranges*** stanza.

   [edit interfaces ge-1/0/0 auto-configure]
   user@host# **edit vlan-ranges**

4. Specify the dynamic VLAN profile that you want the interface to use.

   [edit interfaces ge-1/0/0 auto-configure vlan-ranges]
   user@host# **set dynamic-profile VLAN-PROF1**

Associating a Stacked VLAN Dynamic Profile with an Interface

To associate a stacked (dual-tag) VLAN dynamic profile with an interface:

1. Access the interface that you want to use for creating VLANs.

   [edit interfaces]
   user@host# **edit interfaces ge-1/0/0**

2. Edit the auto-configure stanza to automatically configure the stacked VLANs.

   [edit interfaces ge-1/0/0]
   user@host# **edit** auto-configure

3. Edit the ***stacked-vlan-ranges*** stanza.

   [edit interfaces ge-1/0/0 auto-configure]
   user@host# **edit stacked-vlan-ranges**

4. Specify the dynamic VLAN profile that you want the interface to use.

   [edit interfaces ge-1/0/0 auto-configure stacked-vlan-ranges]
   user@host# **set dynamic-profile STACKED-VLAN-PROF1**

Configuring Which VLAN Ethernet Packet Types Dynamic Profiles Can Accept

To create dynamic single-tag VLANs and dynamic stacked (dual-tag) VLANs, you must specify what Ethernet packet type you want the single-tag VLAN or stacked VLAN dynamic profile to accept. You can configure which VLAN Ethernet packet types a dynamic profile accepts in the following ways:

- Configuring the VLAN Ethernet Packet Type for Single-Tag VLAN Dynamic Profiles on page 721
- Configuring the VLAN Ethernet Packet Type for Stacked VLAN Dynamic Profiles on page 722

Configuring the VLAN Ethernet Packet Type for Single-Tag VLAN Dynamic Profiles

To configure the VLAN Ethernet packet type the VLAN dynamic profile can accept:

1. Access the interface over which you want to create dynamic VLANs.

   user@host# **edit interfaces ge-0/0/0**

2. Edit the VLAN **auto-configure** stanza.

   [edit interfaces ge-0/0/0]
   user@host# **edit** auto-configure

3. Edit the **vlan-ranges** stanza.

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# **edit vlan-ranges**

4. Access the VLAN dynamic profile for which you want to configure VLAN ranges.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **edit dynamic-profile VLAN-PROF1**

5. Specify what VLAN Ethernet packet type the VLAN or stacked VLAN dynamic profile accepts.

   NOTE: This release supports **inet** and **dhcp-v4** for IPv4 packets, **inet6** and **dhcp-v6** for IPv6 packets, and **pppoe** for PPP packets.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges dynamic-profile VLAN-PROF-1]
   user@host# **set accept inet**

Configuring the VLAN Ethernet Packet Type for Stacked VLAN Dynamic Profiles

To configure the VLAN Ethernet packet type the stacked VLAN dynamic profile can accept:

1. Access the interface over which you want to create dynamic VLANs.

   user@host# **edit interfaces ge-0/0/0**

2. Edit the VLAN **auto-configure** stanza.

   [edit interfaces ge-0/0/0]
   user@host# **edit** auto-configure

3. Edit the **stacked-vlan-ranges** stanza.

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# **edit stacked-vlan-ranges**

4. Access the VLAN dynamic profile for which you want to configure VLAN ranges.

   [edit interfaces ge-0/0/0 auto-configure stacked-vlan-ranges]
   user@host# **edit dynamic-profile STACKED-VLAN-PROF1**

5. Specify what VLAN Ethernet packet type the stacked VLAN dynamic profile accepts.

   NOTE: This release supports **inet** and **dhcp-v4** for IPv4 packets, **inet6** and **dhcp-v6** for IPv6 packets, and **pppoe** for PPP packets.

   [edit interfaces ge-0/0/0 auto-configure stacked-vlan-ranges dynamic-profile
   STACKED-VLAN-PROF1]
   user@host# **set accept inet**

Copyright © 2013, Juniper Networks, Inc.

Configuring an Authentication Password for VLAN or Stacked VLAN Ranges

You can specify an authentication password for dynamically created VLAN or stacked VLAN interfaces at the **[edit interfaces *interface-name* auto-configure vlan-ranges authentication]** or **[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication]** hierarchy level. This password is sent to the external AAA authentication server for subscriber authentication.

NOTE: You must configure the username-include statement to enable the use of authentication. The password (Interfaces) statement is not required and does not cause the interface to use authentication if the username-include statement is not included.

To configure an authentication password:

1. Access the interface over which you want to create dynamic VLANs.

   user@host# **edit interfaces ge-0/0/0**

2. Edit the VLAN **auto-configure** stanza.

   [edit interfaces ge-0/0/0]
   user@host# **edit** auto-configure

3. Edit the **vlan-ranges** or **stacked-vlan-ranges** stanza.

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# **edit vlan-ranges**

   or

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# **edit stacked-vlan-ranges**

4. Edit the VLAN **authentication** stanza.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **edit authentication**

5. Specify a password that is sent to the external AAA authentication server for subscriber authentication.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **set password (Interfaces) PSSWD1**

Related Documentation

- Configuring Dynamic Authentication for VLAN Interfaces on page 728

Configuring VLAN Ranges for Use with Dynamic Profiles

You define dynamic VLAN ranges under the **[edit interfaces]** hierarchy. You can configure VLAN ranges in the following ways for use with dynamic profiles:

- Configuring Single-Level VLAN Ranges for Use with VLAN Dynamic Profiles on page 724
- Configuring Stacked VLAN Ranges for Use with Stacked VLAN Dynamic Profiles on page 725
- Configuring Dynamic Mixed VLAN Ranges on page 726
- Configuring VLAN Dynamic Profile Override on page 727

Configuring Single-Level VLAN Ranges for Use with VLAN Dynamic Profiles

You configure VLAN ranges at the **[edit interfaces]** hierarchy level by specifying the **vlan-tagging** statement for the interface and defining up to 32 VLAN ranges for use with a VLAN dynamic profile.

To configure a VLAN range:

1. Access the interface over which you want to create dynamic VLANs.

   user@host# edit interfaces ge-0/0/0

2. Specify the **vlan-tagging** statement to indicate that this interface is for use with stacked VLAN ranges.

   [edit interfaces ge-0/0/0]
   user@host# set vlan-tagging

3. Access the VLAN **[auto-configure]** hierarchy level.

   [edit interfaces ge-0/0/0]
   user@host# edit auto-configure

4. Access the **[vlan-ranges]** hierarchy level.

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# edit vlan-ranges

5. Access the VLAN dynamic profile for which you want to configure VLAN ranges.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# edit dynamic-profile VLAN-PROF1

6. Specify the VLAN ranges that you want the dynamic profile to use. The following example specifies a lower VLAN ID limit of 3000 and any upper VLAN ID limit (a range from 1 through 4094).

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# set ranges 3000-any

NOTE: You can configure multiple VLAN range groups (up to 32 total) on the same physical interface that use different VLAN dynamic profiles.

Copyright © 2013, Juniper Networks, Inc.

7. (Optional) Access another VLAN dynamic profile for which you want to configure VLAN ranges.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **edit dynamic-profile VLAN-PROF2**

8. (Optional) Specify the VLAN ranges that you want the dynamic profile to use. The following example specifies a lower VLAN ID limit of 5000 and any upper VLAN ID limit (a range from 1 through 4094).

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **set ranges 5000-any**

Configuring Stacked VLAN Ranges for Use with Stacked VLAN Dynamic Profiles

You configure stacked VLAN ranges at the **[edit interfaces]** hierarchy level by specifying the **stacked-vlan-tagging** statement for the interface and defining up to 32 stacked VLAN ranges for use with a stacked VLAN dynamic profile.

To configure a VLAN range:

1. Access the interface over which you want to create dynamic VLANs.

   user@host# **edit interfaces ge-0/0/0**

2. Specify the **stacked-vlan-tagging** statement to indicate that this interface is for use with stacked VLAN ranges.

   [edit interfaces ge-0/0/0]
   user@host# **set stacked-vlan-tagging**

3. Access the VLAN **[auto-configure]** hierarchy level.

   [edit interfaces ge-0/0/0]
   user@host# **edit** auto-configure

4. Access the **[stacked-vlan-ranges]** hierarchy level.

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# **edit stacked-vlan-ranges**

5. Access the VLAN dynamic profile for which you want to configure VLAN ranges.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **edit dynamic-profile VLAN-PROF1**

6. Specify the outer and inner stacked VLAN ranges that you want the dynamic profile to use. The following example specifies an outer stacked VLAN ID range from 2000 through 4000 and an inner stacked VLAN ID range of **any** (enabling a range from 1 through 4094 for the inner stacked VLAN ID).

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **set ranges 2000-4000,any**

NOTE: **You can configure multiple dynamic profile associations (up to 32) with different VLAN range groups on each physical interface.**

7. (Optional) Access another VLAN dynamic profile for which you want to configure VLAN ranges.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **edit dynamic-profile VLAN-PROF2**

8. (Optional) Specify the outer and inner stacked VLAN ranges that you want the dynamic profile to use. The following example specifies an outer stacked VLAN ID range from 4001 through 6000 and an inner stacked VLAN ID range of **any** (enabling a range from 1 through 4094 for the inner stacked VLAN ID).

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **set ranges 4001-6000,any**

Configuring Dynamic Mixed VLAN Ranges

Dynamic VLAN and dynamic stacked VLAN configuration supports mixed (or flexible) VLAN ranges. You configure mixed VLAN ranges at the **[edit interfaces]** hierarchy level by specifying the **flexible-vlan-tagging** statement for the interface and defining up to 32 VLAN and stacked VLAN range groups for use with different VLAN or stacked VLAN dynamic profiles.

NOTE: Junos VLAN IDs for single-tag VLANs are equivalent to the outer tags used for stacked (dual-tag) VLANs. When configuring mixed (flexible) VLANs, any overlap on single-tag VLAN IDs and stacked VLAN outer tag values is supported only for dynamic VLANs on MPC line cards. When configuring mixed (flexible) VLANS on DPCE line cards, overlapping single-tag VLAN IDs and stacked VLAN outer tag values is not supported. This means that a dynamically created single-tagged VLAN interface prevents any overlapping stacked VLAN interfaces from being created or a dynamically created stacked VLAN interface prevents any overlapping single-tagged VLAN interfaces from being created.

To configure both VLAN and stacked VLAN ranges:

1. Access the interface over which you want to create dynamic VLANs.

   user@host# **edit interfaces ge-0/0/0**

2. Specify the **flexible-vlan-tagging** statement to indicate that this interface is for use with both VLAN and stacked VLAN ranges.

   [edit interfaces ge-0/0/0]
   user@host# **set flexible-vlan-tagging**

3. Define interface automatic configuration values.

   [edit interfaces ge-0/0/0]
   user@host# **edit** auto-configure

4. Specify that you want to modify VLAN ranges.

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# **edit vlan-ranges**

Copyright © 2013, Juniper Networks, Inc.

5. Access the VLAN dynamic profile for which you want to configure VLAN ranges.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **edit dynamic-profile VLAN-PROF1**

6. Specify the VLAN ranges that you want the dynamic profile to use. The following example specifies a lower VLAN ID limit of 2000 and an upper VLAN ID limit of 3000.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **set ranges 2000-3000**

   NOTE: **You can configure multiple dynamic profile associations (up to 32) with different VLAN range groups on each physical interface.**

7. Specify that you want to modify stacked VLAN ranges.

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# **edit stacked-vlan-ranges**

8. Access the VLAN dynamic profile for which you want to configure VLAN ranges.

   [edit interfaces ge-0/0/0 auto-configure stacked-vlan-ranges]
   user@host# **edit dynamic-profile VLAN-PROF2**

9. Specify the outer and inner stacked VLAN ranges that you want the dynamic profile to use. The following example specifies an outer stacked VLAN ID range from 3001 through 4000 (to avoid overlapping VLAN IDs with single-tag VLANs) and an inner stacked VLAN ID range of **any** (enabling a range from 1 through 4094 for the inner stacked VLAN ID).

   [edit interfaces ge-0/0/0 auto-configure stacked-vlan-ranges]
   user@host# **set ranges 3001-4000,any**

   NOTE: **You can configure multiple dynamic profile associations (up to 32) with different VLAN range groups on each physical interface.**

Configuring VLAN Dynamic Profile Override

You can override dynamic profile assignment to individual VLANs that are already part of a previously defined VLAN range. This functionality provides a type of exception to an assigned VLAN range. It enables you to configure individual VLAN IDs to use a different dynamic profile from the one assigned to the VLAN range that includes the individual VLAN ID.

To configure dynamic profile override for a specific VLAN:

1. Access the interface on which you want to create a dynamic profile override.

   user@host# **edit interfaces ge-0/0/0**

2. Access the interface automatic configuration hierarchy.

   [edit interfaces ge-0/0/0]
   user@host# **edit** auto-configure

3. Access either the single-tagged or dual-tagged (stacked) VLAN ranges that you want to modify.

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# **edit vlan-ranges**

   or

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# **edit stacked-vlan-ranges**

4. Define the **override** statement along with the VLAN tag that you want to override and the dynamic profile that you want to use when overriding the specified VLAN tag.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# **set override tag 20 dynamic-profile NewProfile**

   or

   [edit interfaces ge-0/0/0 auto-configure stacked-vlan-ranges]
   user@host# **set override tag 20 dynamic-profile NewProfile**

Configuring Dynamic Authentication for VLAN Interfaces

You can use dynamic profiles, in conjunction with RADIUS, to dynamically create logical VLAN interfaces in the default logical system and in a specified routing instance. As DHCP clients in the same VLAN become active, corresponding interfaces are assigned to any specified routing instances. You can also dynamically create an underlying VLAN interface for incoming subscribers, associate interfaces created on this VLAN with the default logical system and a specified routing instance, and define RADIUS authentication values for the dynamically created interfaces.

Before you configure dynamic VLAN authentication, configure DHCP Local Server or DHCP Relay over which you want the dynamic VLAN interfaces to function.

For information about DHCP Local Server or DHCP Relay, see:

- Extended DHCP Local Server Overview on page 200
- Extended DHCP Relay Agent Overview on page 276

NOTE: You can also configure dynamically created VLAN interfaces over PPP or PPPoE interfaces. For information about how to configure PPP or PPPoE, see the *Junos OS Network Interfaces Library for Routing Devices*.

Copyright © 2013, Juniper Networks, Inc.

To configure dynamic authentication for dynamically created VLAN interfaces:

1. Configure an access profile that contains the appropriate accounting order, authentication order, and server access values.

   For information about how to configure an access profile, RADIUS accounting, RADIUS statistics, and how to define RADIUS server access, see:

   - Configuring an Access Profile for Subscriber Management on page 149
   - Specifying the Authentication and Accounting Methods for Subscriber Access on page 24
   - Configuring Per-Subscriber Session Accounting on page 29
   - Configuring Router or Switch Interaction with RADIUS Servers on page 23

2. Configure a dynamic profile that uses the default logical system and creates specific routing instances to contain dynamically created VLAN interfaces.

   See "Configuring a VLAN Dynamic Profile That Associates VLAN Interfaces with Separate Routing Instances" on page 718.

3. Define the VLAN physical interface for automatic configuration.

   See the following topics:

   - *Enabling VLAN Tagging*
   - Configuring Which VLAN Ethernet Packet Types Dynamic Profiles Can Accept on page 721
   - Configuring VLAN Ranges for Use with Dynamic Profiles on page 724
   - Configuring an Authentication Password for VLAN or Stacked VLAN Ranges on page 723
   - Configuring VLAN Interface Username Information for AAA Authentication on page 729

4. Associate an access profile to the VLAN interface.

   See "Attaching Access Profiles" on page 150.

5. Associate a dynamic profile to the VLAN interface.

   See "Configuring VLAN Interfaces to Use Dynamic Profiles" on page 720.

Related Documentation

- Dynamic 802.1Q VLAN Overview on page 701

Configuring VLAN Interface Username Information for AAA Authentication

You can define interface information that is included in the username that is subsequently passed to the external AAA authentication service (for example, RADIUS) when creating dynamic VLANs or stacked VLANs. The AAA authentication service uses this information to authenticate the VLAN or stacked VLAN physical interface. After the interface is

authenticated, the AAA service can send the required routing instance values to the system for use in dynamically creating VLAN or stacked VLAN interfaces.

NOTE: The following example configures username information on VLANs. However, you can also configure dynamic authentication on stacked VLANs by configuring the same statements at the **[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication]** hierarchy level.

To configure VLAN interface username information:

1. Access the interface over which you want to configure username information.

   user@host# edit interfaces ge-0/0/0

2. Edit the auto-configure stanza.

   [edit interfaces ge-0/0/0]
   user@host# edit auto-configure

3. Edit the ***vlan-ranges*** stanza.

   [edit interfaces ge-0/0/0 auto-configure]
   user@host# edit vlan-ranges

4. Edit the ***authentication*** stanza.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# edit authentication

5. Edit the username-include stanza.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges]
   user@host# edit username-include

6. Specify the username statements that you want the AAA authentication service to use to authenticate the username.

   [edit interfaces ge-0/0/0 auto-configure vlan-ranges authentication username-include]
   user@host# set delimiter

Related Documentation

- Configuring Dynamic Authentication for VLAN Interfaces on page 728
- Option 82 Suboptions in Authentication Usernames for Autosense VLANs on page 730

Option 82 Suboptions in Authentication Usernames for Autosense VLANs

You can specify the Option 82 suboptions that are concatenated with the username during the authentication process for autosense VLANs. The option 82 value used in creating the username is based on the option 82 value that is encoded in the incoming DHCP discover packet.

You can specify either, both, or neither of the Agent Circuit ID (suboption 1) and the Agent Remote ID (suboption 2). 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

 Copyright © 2013, Juniper Networks, Inc.

supplied, the raw payload of Option 82 from the PDU is concatenated to the username. The use of Option 82 suboptions is supported for DHCPv4 discover packets only.

Related Documentation

- Configuring VLAN Interface Username Information for AAA Authentication on page 729

Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs

For DHCPv4, Option 82 has suboptions containing the ACI and ARI that are concatenated with the username during the authentication process for autosense (dynamic) VLANs. For DHCPv6, the relay agent uses Options 18 and Option 37 to convey the ACI and ARI, respectively. You can include these options in the username to generate unique usernames that identify subscribers for authentication in DHCPv6 dynamic VLANs.

A DHCPv6 Solicit message encapsulated with a Relay-Forward message header and one without the Relay-Forward message header are eligible for dynamic VLAN creation when you configure the DHCPv6 packet type for autosensing. Options 18 and Option 37 are provided in the Relay-Forward message header and are extracted only from this header and not from the options within the DHCPv6 Solicit message. In addition, if the DHCPv6 Solicit message is encapsulated in multiple Relay-Forward message headers, only the option values from the innermost Relay-Forward message header are used for username authentication. If these options are sent by the client or DHCPv6 relay agent, and if dynamic VLAN authentication is configured to use these options in the username, then the options are included in the username for authentication. If either of these options is not sent by the client or DHCPv6 relay agent, or if the dynamic VLAN authentication is not configured to use the option in the username, the username is constructed without the option.

To include Option 18 or Option 37 in the username for DHCPv6 dynamic VLANs, include the **option-18** and **option-37** statements at the **[edit interfaces *interface-name* auto-configure vlan-ranges authentication** username-include**]** hierarchy level. To include Options 18 or Option 37 in the username for stacked VLANs, include **option-18** and **option-37** statements at the **[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication** username-include**]** hierarchy level.

Related Documentation

- Configuring VLAN Interface Username Information for AAA Authentication on page 729
- **username-include on page 2098**
- **option-18 on page 1824**
- **option-37 on page 1825**
- Configuring Which VLAN Ethernet Packet Types Dynamic Profiles Can Accept on page 721

Automatically Removing VLANs with No Subscribers

You can always clear or delete subscriber VLANs manually. However, you can also configure the interface to automatically remove dynamic subscriber VLANs when no client sessions (for example, DHCP or PPPoE) exist on the VLAN.

When configuring automatic removal of dynamic subscriber VLANs, keep the following in mind:

- You can configure automatic VLAN removal only on individual physical interfaces. You cannot configure the feature globally.
- Automatic VLAN removal is not supported for use on Layer 2 Wholesale interfaces.
- PPPoE subscriber interfaces require the use of a dynamic profiles when configured over dynamic VLANS. However, dynamic profiles are not required for use with DHCP subscriber interfaces that use underlying dynamic VLANs. Because the remove-when-no-subscribers functionality triggers when no dynamic client sessions exist on a dynamic VLAN, automatic removal of underlying dynamic VLANs is not supported when DHCP subscriber interfaces are not created using dynamic profiles.
- The **maintain-subscriber** statement and **remove-when-no-subscribers** statement are mutually exclusive. When the router is configured to maintain subscribers, you cannot also specify that dynamically configured VLAN interfaces are removed when no subscribers exist.
- If PPPoE subscriber session lockout is also configured, the router does not remove the unused subscriber VLAN until the lockout time has expired for each client undergoing lockout on the underlying interface.

To configure automatic removal of subscriber VLANs when no client sessions exist on the VLAN:

1. Access the interface for which you want to enable automatic removal of subscriber VLANs.

   user@host# edit interfaces ge-1/1/1

2. Access the interface automatic configuration hierarchy.

   [edit interfaces ge-1/1/1]
   user@host# edit auto-configure

3. Enable subscriber VLAN removal with the **remove-when-no-subscribers** statement.

   [edit interfaces ge1/1/1 auto-configure]
   user@host# set remove-when-no-subscribers

Related Documentation

- Dynamic 802.1Q VLAN Overview on page 701
- Configuring VLAN Interfaces to Use Dynamic Profiles on page 720
- PPPoE Subscriber Session Lockout Overview on page 893

Copyright © 2013, Juniper Networks, Inc.

Configuring Dynamic VLANs Based on Agent Circuit Identifier Information

On MX Series routers with Modular Port Concentrators/Modular Interface Cards (MPCs/MICs) that face the access side of the network, you can configure dynamic VLAN subscriber interfaces based on agent circuit identifier (ACI) information, also known as *ACI-based dynamic VLANs*, for DHCP and PPPoE subscribers. To do so, you create an *ACI interface set*, which is a logical collection of subscriber interfaces that originate at the same household or on the same access-loop port, and then reference the ACI interface set in the dynamic profile for a PPPoE or IP demultiplexing (IP demux) logical subscriber interface.

Grouping subscriber interfaces into ACI interface sets to create ACI-based dynamic VLANs facilitates application of subscriber-based services, such as class of service (CoS) and interface-shared filters, to all of the subscriber interfaces from the same household.

Before you begin:

1. Configure the underlying physical interface for single-tag VLANs or stacked (dual-tag) VLANs.

   See the following topics:

   - *802.1Q VLANs Overview*
   - Configuring VLAN Dynamic Profiles on page 714
   - Configuring VLAN Interfaces to Use Dynamic Profiles on page 720

2. Create a dynamic profile that defines the logical subscriber interface.

   See the following topics:

   - Configuring a Basic Dynamic Profile on page 673
   - Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles on page 899
   - Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771

To configure a dynamic VLAN subscriber interface based on ACI information:

1. Create a dynamic profile that defines the dynamic ACI interface set.

   See "Defining Agent Circuit Identifier Interface Sets" on page 735.

2. (Optional) Include attributes for PPPoE, CoS, and interface-shared filters in the dynamic profile for the ACI interface set.

   See "Defining Agent Circuit Identifier Interface Sets" on page 735.

3. (Optional) In the dynamic profile for the ACI interface set, configure the router to use the Actual-Data-Rate-Downstream VSA [26-130] or Access-Loop-Encapsulation VSA [26-144] value in PPPoE control packets to adjust CoS shaping-rate and overhead-accounting attributes at a per-household level.

Copyright © 2013, Juniper Networks, Inc.

See "Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs" on page 742.

4. Dynamically or statically configure the underlying VLAN logical interface to enable dynamic subscriber interface creation based on ACI information.

   - For dynamic underlying VLAN interfaces, see "Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information" on page 737.
   - For static underlying VLAN interfaces, see "Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information" on page 739.

5. Associate the dynamic ACI interface set with the dynamic PPPoE or dynamic IP demux logical subscriber interface.

   See "Configuring Dynamic VLAN Subscriber Interfaces Based on Agent Circuit Identifier Information" on page 740.

6. (Optional) In the dynamic profile for the PPPoE (**pp0**) subscriber interface, configure the router to use the Actual-Data-Rate-Downstream VSA [26-130] or Access-Loop-Encapsulation VSA [26-144] value in PPPoE control packets to adjust CoS shaping-rate and overhead-accounting attributes at a per-subscriber level.

   See "Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs" on page 742.

7. (Optional) Verify the ACI-based dynamic VLAN subscriber interface configuration.

   See "Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration" on page 747.

8. (Optional) Clear the ACI interface set from the router when the interface set no longer has any active subscriber sessions.

   See "Clearing Agent Circuit Identifier Interface Sets" on page 748.

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Overview on page 702
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704
- Agent Circuit Identifier-Based Dynamic VLANs Bandwidth Management Overview on page 706
- Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration on page 747
- Clearing Agent Circuit Identifier Interface Sets on page 748

Copyright © 2013, Juniper Networks, Inc.

Defining Agent Circuit Identifier Interface Sets

To configure the router to create dynamic VLAN subscriber interfaces for DHCP and PPPoE subscribers based on agent circuit identifier (ACI) information, you must create a dynamic ACI interface set. An *ACI interface set* is a logical collection of subscriber interfaces that originate at the same household or on the same access-loop port.

Because DHCP and PPPoE traffic sent to the router from the same household carries the same ACI value in DHCP and PPPoE control packets, the router groups these subscriber interfaces into a single ACI interface set. Grouping subscriber interfaces into ACI interface sets facilitates application of attributes for PPPoE, class of service (CoS), and interface-shared filters to all of the subscriber interfaces from the same household.

To configure an ACI interface set in a dynamic profile:

1. Name the dynamic profile that defines the ACI interface set.

   [edit]
   user@host# **edit** dynamic-profiles ***profile-name***

2. Specify that you want to configure the interfaces for the dynamic profile.

   [edit dynamic-profiles *profile-name*]
   user@host# **edit** interfaces

3. Configure the dynamic ACI interface set.

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

   You must use the **$junos-interface-set-name** predefined dynamic variable to represent the name of the ACI interface set. The **$junos-interface-set-name** variable is dynamically replaced with the actual ACI interface set name generated by the router when the first DHCP or PPPoE subscriber from that household logs in.

4. Include the interfaces for the dynamic ACI interface set.

   [edit dynamic-profiles *profile-name* interfaces interface-set
   "$junos-interface-set-name"]
   user@host# **set** interface **$junos-interface-ifd-name**

   You must use the **$junos-interface-ifd-name** predefined dynamic variable to represent the name of the ACI interface set. The **$junos-interface-ifd-name** variable is dynamically replaced with the name of the interface on which the DHCP or PPPoE subscriber accesses the router.

   NOTE: The unit *logical-unit-number* statement is not required in the dynamic profile at the [edit dynamic-profiles *profile-name* interfaces interface-set *interface-set-name* interface *interface-name*] hierarchy level when you configure an ACI interface set.

5. (Optional) For dynamic PPPoE subscriber interfaces, configure the maximum number of dynamic PPPoE sessions that the router can activate for the ACI interface set.

```
[edit dynamic-profiles profile-name interfaces interface-set
  "$junos-interface-set-name"]
user@host# edit pppoe-underlying-options
[edit dynamic-profiles profile-name interfaces interface-set "$junos-interface-set-name"
  pppoe-underlying-options]
user@host# set max-sessions number
```

Issuing the **max-sessions** statement in a dynamic profile for an ACI interface set limits the maximum number of dynamic PPPoE sessions at the ACI interface set level from the same household.

6. (Optional) Apply attributes for CoS and interface filters to all subscriber interfaces belonging to the ACI interface set.

The following example shows the minimum dynamic profile required to define an ACI interface set named aci-vlan-set-profile. The **interface-set** stanza uses predefined dynamic variables to represent the interface set (**$junos-interface-set-name**) and the underlying physical interface (**$junos-interface-ifd-name**).

```
[edit dynamic-profiles aci-vlan-set-profile]
interfaces {
  interface-set "$junos-interface-set-name" {
    interface "$junos-interface-ifd-name";
  }
}
```

The following example shows a more complex dynamic profile for an ACI interface set named aci-vlan-set-profile-pppoe-cos. In addition to the required **interface-set** stanza, this profile includes optional attributes for PPPoE (**max-sessions** statement) and CoS. The router applies these PPPoE and CoS attributes to all subscriber interfaces from the same household, which is represented by the ACI interface set.

```
[edit dynamic-profiles aci-vlan-set-profile-pppoe-cos]
variables {
  ds1q1q2DP uid;
  ef1_dp uid;
}
  interfaces {
    interface-set "$junos-interface-set-name" {
      interface "$junos-interface-ifd-name";
      pppoe-underlying-options {
        max-sessions 3;
      }
    }
  }
  class-of-service {
    traffic-control-profiles {
      tcp2 {
        scheduler-map "$ds1q1q2DP";
        shaping-rate 50m;
        overhead-accounting bytes -20;
        guaranteed-rate 30m;
      }
    }
    interfaces {
```

 Copyright © 2013, Juniper Networks, Inc.

```
        interface-set "$junos-interface-set-name" {
            output-traffic-control-profile tcp2;
        }
    }
    scheduler-maps {
        "$ds1q1q2DP" {
            forwarding-class ef scheduler "$ef1_dp";
        }
    }
    schedulers {
        "$ef1_dp" {
            transmit-rate percent 25;
            priority low;
        }
    }
}
```

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Overview on page 702
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration on page 747
- Clearing Agent Circuit Identifier Interface Sets on page 748
- Applying CoS Attributes to VLANs Using Agent-Circuit-Identifiers on page 984
- Example: Interface-Shared Filter Configuration on page 1217

Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information

After you define the agent circuit identifier (ACI) interface set, you must configure the underlying VLAN interface to enable creation of dynamic VLAN subscriber interfaces based on ACI information. You can configure the underlying VLAN interface statically or dynamically.

This topic describes how to configure the underlying VLAN interface *dynamically*.

Before you begin:

- Create a dynamic profile that defines the underlying VLAN interface.

  See the following topics:

  - Configuring a Basic Dynamic Profile on page 673
  - Configuring VLAN Dynamic Profiles on page 714

- Configuring VLAN Interfaces to Use Dynamic Profiles on page 720

To configure a dynamic underlying VLAN interface to use ACI information:

- In the dynamic profile for the underlying VLAN interface, associate the dynamic profile that defines the ACI interface set with the underlying VLAN interface.

  [edit dynamic-profiles *profile-name*]
  user@host# set interfaces *interface-name* unit *logical-unit-number* auto-configure agent-circuit-identifier dynamic-profile *aci-interface-set-profile-name*

  For example, the following statement in a dynamic profile named aci-vlan-underlying-profile-demux associates the dynamic underlying VLAN interface with dynamic profile aci-vlan-set-profile2 that defines the ACI interface set. You must use the predefined dynamic variable **$junos-interface-ifd-name** to represent the interface name, and **$junos-interface-unit** to represent the logical unit number.

  [edit dynamic-profiles aci-vlan-underlying-profile-demux]
  user@host# set interfaces “$junos-interface-ifd-name” unit “$junos-interface-unit” auto-configure agent-circuit-identifier dynamic-profile aci-vlan-set-profile2

The following example shows the dynamic configuration that uses this statement. This configuration enables the underlying dynamic IP demultiplexing (IP demux) VLAN interface to create dynamic subscriber interfaces based on ACI information by applying a single default ACI interface set dynamic profile (aci-vlan-set-profile2) to all households on the VLAN interface.

```
[edit dynamic-profiles aci-vlan-underlying-profile-demux]
interfaces {
  “$junos-interface-ifd-name” {
    unit "$junos-interface-unit" {
      auto-configure {
        agent-circuit-identifier {
          dynamic-profile aci-vlan-set-profile2;
        }
      }
      vlan-id "$junos-vlan-id";
      demux-options {
        underlying-interface “$junos-interface-ifd-name”;
      }
      family inet {
        unnumbered-address lo0.0 preferred-source-address 100.20.0.2;
      }
    }
  }
}
```

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Overview on page 702
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration on page 747

Copyright © 2013, Juniper Networks, Inc.

Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information

After you define the agent circuit identifier (ACI) interface set, you must configure the underlying VLAN interface to enable creation of dynamic VLAN subscriber interfaces based on ACI information. You can configure the underlying VLAN interface statically or dynamically.

This topic describes how to configure the underlying VLAN interface *statically*.

To configure a static underlying VLAN interface to use ACI information:

- Associate the dynamic profile that defines the ACI interface set with the static underlying VLAN interface.

  ```
  [edit]
  user@host# set interfaces interface-name unit logical-unit-number auto-configure
    agent-circuit-identifier dynamic-profile aci-interface-set-profile-name
  ```

  For example, the following statement associates static Gigabit Ethernet VLAN interface ge-1/0/0.0 with the dynamic profile aci-vlan-set-profile that defines the ACI interface set.

  ```
  [edit]
  user@host# set interfaces ge-1/0/0 unit 0 auto-configure agent-circuit-identifier
    dynamic-profile aci-vlan-set-profile
  ```

The following example shows the static configuration that uses this statement. This configuration enables the underlying VLAN interface ge-1/0/0.0 to create dynamic subscriber interfaces based on ACI information by applying a single default ACI interface set dynamic profile (aci-vlan-set-profile) to all households on the VLAN interface.

```
[edit]
interfaces {
  ge-1/0/0 {
    flexible-vlan-tagging;
    unit 0 {
      vlan-id 100;
      auto-configure {
        agent-circuit-identifier {
          dynamic-profile aci-vlan-set-profile;
        }
      }
    }
  }
}
```

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Overview on page 702
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration on page 747

Copyright © 2013, Juniper Networks, Inc.

Configuring Dynamic VLAN Subscriber Interfaces Based on Agent Circuit Identifier Information

After you define the dynamic agent circuit identifier (ACI) interface set and enable creation of ACI-based dynamic VLAN subscriber interfaces on the underlying VLAN interface, you must complete the configuration by associating the ACI interface set with the PPPoE or IP demultiplexing (IP demux) subscriber interface in the dynamic profile for the subscriber interface.

Before you begin:

- Create a dynamic profile that defines the logical subscriber interface.

  See the following topics:

  - Configuring a Basic Dynamic Profile on page 673
  - Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles on page 899
  - Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771

To configure a dynamic VLAN subscriber interface based on ACI information:

- In the dynamic profile for the PPPoE or IP demux subscriber interface, associate the dynamic ACI interface set with the dynamic VLAN subscriber interface name (**pp0** or **demux0**) and logical unit number.

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

  For example, the following statement in a dynamic profile named aci-vlan-pppoe-profile associates the dynamic ACI interface set with the dynamic **pp0** (PPPoE) logical subscriber interface. You must use the predefined dynamic variable **$junos-interface-set-name** to represent the name of the dynamic ACI interface set, and **$junos-interface-unit** to represent the logical unit number of the subscriber interface.

  ```
  [edit dynamic-profiles aci-vlan-pppoe-profile]
  user@host# set interfaces interface-set $junos-interface-set-name interface pp0 unit
    $junos-interface-unit
  ```

  Similarly, the following statement in a dynamic profile named aci-vlan-demux-profile associates the dynamic ACI interface set (represented by **$junos-interface-set-name**) with the **demux0** (IP demux) logical subscriber interface.

  ```
  [edit dynamic-profiles aci-vlan-demux-profile]
  user@host# set interfaces interface-set $junos-interface-set-name interface demux0
    unit $junos-interface-unit
  ```

The following examples show the dynamic configurations that use each of these statements. The following sample configuration shows a dynamic profile named aci-vlan-pppoe-profile for an ACI-based dynamic PPPoE (**pp0**) subscriber interface for use by PPPoE subscribers.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit dynamic-profiles aci-vlan-pppoe-profile]
interfaces {
  interface-set "$junos-interface-set-name" {
    interface pp0 {
      unit "$junos-interface-unit";
    }
  }
  pp0 {
    unit "$junos-interface-unit" {
      ppp-options {
        chap;
        pap;
      }
      pppoe-options {
        underlying-interface "$junos-underlying-interface";
        server;
      }
      no-keepalives;
      family inet {
        unnumbered-address lo0.0;
      }
    }
  }
}
```

The following sample configuration shows a dynamic profile named aci-vlan-demux-profile for an ACI-based dynamic IP demux(**demux0**) subscriber interface for use by DHCP subscribers.

```
[edit dynamic-profiles aci-vlan-demux-profile]
interfaces {
  interface-set "$junos-interface-set-name" {
    interface demux0 {
      unit "$junos-interface-unit";
    }
  }
  demux0 {
    unit "$junos-interface-unit" {
      demux-options {
        underlying-interface "$junos-underlying-interface";
      }
      family inet {
        demux-source {
          $junos-subscriber-ip-address;
        }
        unnumbered-address lo0.0 preferred-source-address 100.20.200.202;
      }
    }
  }
}
```

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Overview on page 702
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733

- Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration on page 747
- Clearing Agent Circuit Identifier Interface Sets on page 748

Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs

You can configure the router to use either or both of the Actual-Data-Rate-Downstream [26-130] or Access-Loop-Encapsulation [26-144] DSL Forum vendor-specific attribute (VSA) values in PPPoE control packets to adjust the CoS shaping-rate and overhead-accounting attributes, respectively, for dynamic agent circuit identifier (ACI) interface sets and their associated ACI-based dynamic VLAN subscriber interfaces.

Before you begin:

- To configure adjustment of the CoS shaping rate and overhead accounting attributes on a per-household basis, create a dynamic profile that defines the dynamic ACI interface set.

  See "Defining Agent Circuit Identifier Interface Sets" on page 735.

- To configure adjustment of the CoS shaping rate and overhead accounting attributes on a per-subscriber basis, create a dynamic profile that defines the ACI-based dynamic PPPoE (**pp0**) subscriber interface associated with the ACI interface set.

  See "Configuring Dynamic VLAN Subscriber Interfaces Based on Agent Circuit Identifier Information" on page 740.

To configure the router to use the Actual-Data-Rate-Downstream or Access-Loop-Encapsulation VSA values in PPPoE control packets to adjust the CoS shaping-rate and overhead-accounting attributes for dynamic ACI interface sets and associated ACI-based dynamic VLAN subscriber interfaces, do either or both of the following:

- In a dynamic profile for an ACI interface set or a dynamic profile for an ACI-based PPPoE subscriber interface, configure adjustment of the CoS shaping-rate attribute based on the value of the Actual-Data-Rate-Downstream VSA.

  [edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]
  user@host# **set vendor-specific-tags actual-data-rate-downstream**

- In a dynamic profile for an ACI interface set or a dynamic profile for an ACI-based PPPoE subscriber interface, configure adjustment of the CoS overhead-accounting attribute based on the value of the Access-Loop-Encapsulation VSA.

  [edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]
  user@host# **set vendor-specific-tags access-loop-encapsulation**

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Bandwidth Management Overview on page 706

Copyright © 2013, Juniper Networks, Inc.

- Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ACI Interface Sets on page 709
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733

Configuring Ethernet OAM Support for Service VLANs with Double-Tagged Customer VLANs

On MX Series routers with MPC/MIC interfaces, you can enable propagation of the Ethernet IEEE 802.1ag Operation, Administration, and Maintenance (OAM) state of a static single-tagged service VLAN (S-VLAN) to the dynamic or static double-tagged customer VLAN (C-VLAN) that has the same S-VLAN (outer) tag as the S-VLAN, and, by extension, to subscriber interfaces configured on the C-VLAN. The static S-VLAN logical interface must be configured on a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet physical interface.

Before you begin:

- Make sure the static single-tagged S-VLAN logical interface is configured with the Ethernet 802.1ag OAM connectivity fault management (CFM) continuity check protocol.

  See *IEEE 802.1ag OAM Connectivity Fault Management Overview*.

To enable propagation of the Ethernet OAM state of a static single-tagged S-VLAN to dynamic or static double-tagged C-VLAN logical interfaces:

- Configure a Gigabit Ethernet (ge), 10-Gigabit Ethernet (xe), or aggregated Ethernet (ae) physical interface to propagate the S-VLAN Ethernet OAM state to C-VLAN logical interfaces that have the same S-VLAN (outer) tag as the S-VLAN interface.

  ```
  [edit]
  user@host# set interfaces interface-name-fpc/pic/port oam-on-svlan
  ```

  For example, the following statement enables propagation of the Ethernet OAM state of a static single-tagged S-VLAN on Gigabit Ethernet interface ge-1/0/5 to a dynamic or static double-tagged C-VLAN logical interface with the same S-VLAN (outer) tag as the S-VLAN interface.

  ```
  [edit]
  user@host# set interfaces ge-1/0/5 oam-on-svlan
  ```

  Including the **oam-on-svlan** statement when you configure a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet physical interface causes the router to bring down both of the following when the CFM continuity check protocol detects that the OAM state of the S-VLAN logical interface is down:

  - All dynamic or static double-tagged C-VLANs on the S-VLAN interface that have the same S-VLAN (outer) tag as the S-VLAN interface.
  - All DHCP, IP demultiplexing (IP demux), and PPPoE logical subscriber interfaces configured on the associated C-VLANs.

Example: Gigabit Ethernet Interface with Static S-VLAN, Dynamic C-VLAN, and Dynamic PPPoE Subscriber Interfaces

The following example shows a dynamic subscriber access configuration that uses the **oam-on-svlan** statement on a Gigabit Ethernet interface. This example configures Gigabit Ethernet physical interface ge-1/0/5 with a static single-tagged S-VLAN logical interface (ge-1/0/5.1) that runs the Ethernet 802.1ag OAM CFM continuity check protocol. A dynamic profile named double-vlans creates a dynamic double-tagged C-VLAN interface, and a dynamic profile named pppoe-profile creates dynamic PPPoE subscriber interfaces on the C-VLAN interface. The **oam-on-svlan** statement for ge-1/0/5 propagates the Ethernet OAM state of S-VLAN interface ge-1/0/5.1 to the C-VLAN interface and the dynamic PPPoE subscriber interfaces.

For clarity, the configuration is divided into five steps.

1. Configure a dynamic profile named double-vlans that defines a dynamic double-tagged C-VLAN logical interface.

```
[edit]
dynamic-profiles {
  double-vlans {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
          vlan-tags outer "$junos-stacked-vlan-id" inner "$junos-vlan-id";
          encapsulation ppp-over-ether;
          pppoe-underlying-options {
            dynamic-profile pppoe-profile;
          }
        }
      }
    }
  }
}
```

2. Configure a dynamic profile named pppoe-profile that defines dynamic PPPoE subscriber interfaces on the C-VLAN.

```
[edit]
dynamic-profiles {
  pppoe-profile {
    interfaces {
      pp0 {
        unit "$junos-interface-unit" {
          pppoe-options {
            underlying-interface "$junos-underlying-interface";
            server;
          }
          family inet {
            unnumbered-address lo0.0;
          }
```

 Copyright © 2013, Juniper Networks, Inc.

```
            }
          }
        }
      }
    }
```

3. Configure Gigabit Ethernet physical interface ge-1/0/5.

```
[edit]
interfaces {
  ge-1/0/5 {
    description "connect to remote router";
    flexible-vlan-tagging;
    oam-on-svlan;
    unit 1 {
      vlan-id 1;
    }
    auto-configure {
      stacked-vlan-ranges {
        dynamic-profile double-vlans {
          accept any;
          ranges {
            any,any;
          }
        }
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 100.1.1.1/32 {
          primary;
        }
      }
    }
  }
}
```

The preceding example in Step 3 configures a static, single-tagged S-VLAN logical interface (ge-1/0/5.1) with VLAN ID 1, and references the double-vlans dynamic profile to create a dynamic double-tagged C-VLAN logical interface with S-VLAN (outer) tag **any** and C-VLAN (inner) tag **any**. The tag value **any** represents the entire range of VLAN IDs or S-VLAN IDs, including VLAN ID 1.

Because the C-VLAN outer tag (**any**) matches the S-VLAN tag VLAN ID 1, the **oam-on-svlan** statement in the configuration causes the router to propagate the Ethernet OAM state of S-VLAN ge-1/0/5.1 to the dynamic double-tagged C-VLAN logical interface (created by the double-vlans dynamic profile) and, by extension, to the dynamic PPPoE subscriber interfaces on the C-VLAN (created by the pppoe-profile dynamic profile).

4. Configure the Ethernet 802.1ag OAM CFM continuity check protocol on the static S-VLAN interface (ge-1/0/5.1).

```
[edit]
```

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        action-profile myDefault {
          default-actions {
            interface-down;
          }
        }
        maintenance-domain md1 {
          level 1;
          maintenance-association ma1 {
            continuity-check {
              interval 1s;
            }
            mep 100 {
              interface ge-1/0/5.1;
              direction down;
              remote-mep 101 {
                action-profile myDefault;
              }
            }
          }
        }
      }
    }
  }
}
```

If the CFM continuity check protocol detects that the Ethernet OAM state of S-VLAN interface ge-1/0/5.1 is down, the **interface-down** action in the myDefault action profile causes the router to bring down both of the following:

- The dynamic double-tagged C-VLAN logical interface that has the same S-VLAN (outer) tag as S-VLAN interface ge-1/0/5.1
- The dynamic PPPoE subscriber interfaces configured on the dynamic C-VLAN interface

5. Create a PPP access profile.

   For brevity, this configuration is only partially shown. The missing portions of the configuration are replaced with ellipses (...).

```
[edit]
access {
  ...
  profile ppp-authenticator {
    ...
  }
}
```

Related Documentation

- Ethernet OAM Support for Service VLANs Overview on page 709
- *IEEE 802.1ag OAM Connectivity Fault Management Overview*

Copyright © 2013, Juniper Networks, Inc.

Verifying and Managing Dynamic VLAN Configuration

Purpose View or clear information about dynamic VLANs and stacked VLANs.

Action

- To display subscriber dynamic VLAN information:

  user@host>**show subscribers detail**

- To display interface-specific output for dynamic VLANs:

  user@host>**show interfaces** ***interface-name***

- To clear the binding state of dynamic VLAN interfaces:

  user@host> **clear auto-configuration interfaces**

Related Documentation

- *Junos OS Operational Mode Commands*
- *Junos OS Operational Mode Commands*

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

Purpose View information about dynamic agent circuit identifier (ACI) interface sets and ACI-based dynamic VLAN subscriber interfaces configured on the router.

Action

- To display the logical and physical interface associations for the classifier, rewrite rules, scheduler map objects, and CoS adjustment settings:

  `user@host>` **show class-of-service interface** ***interface-name***

- To display the CoS associations for the specified dynamic ACI interface set:

  `user@host>` **show class-of-service interface-set** ***aci-interface-set-name***

- To display information about the specified CoS traffic shaping and scheduling profile:

  `user@host>` **show class-of-service traffic-control-profile** ***profile-name***

- To display address bindings and ACI interface set information in the client table on the extended DHCP local server:

  `user@host>` **show dhcp server binding detail**

- To display status information about a specified Gigabit Ethernet interface:

  `user@host>` **show interfaces ge-*fpc/pic/port*.*logical-unit-number***

- To display status information about a specified IP demultiplexing (IP demux) interface:

  `user@host>` **show interfaces demux0.*logical-interface-number***

- To display information about all dynamic ACI interface sets configured on the router:

  `user@host>` **show interfaces interface-set**

- To display session-specific information about ACI-based dynamic PPPoE subscriber interfaces:

  `user@host>` **show pppoe interfaces pp0.*logical-unit-number***

- To display information about PPPoE underlying interfaces, including whether creation of ACI-based dynamic VLAN subscriber interfaces is enabled on the underlying interface:

```
user@host> show pppoe underlying-interfaces logical-interface-name detail
```

- To display information about active subscriber sessions associated with ACI interface sets:

```
user@host> show subscribers detail
```

- To display information about active subscriber sessions associated with a specified ACI interface set:

```
user@host> show subscribers aci-interface-set-name aci-interface-set-name detail
```

- To display information about active subscriber sessions that have an agent circuit identifier value containing a matching substring:

```
user@host> show subscribers agent-circuit-identifier agent-circuit-identifier-substring detail
```

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Overview on page 702
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Clearing Agent Circuit Identifier Interface Sets on page 748
- *Junos OS Operational Mode Commands*

Clearing Agent Circuit Identifier Interface Sets

Purpose Clear a specified dynamic agent circuit identifier (ACI) interface set configured on the router.

Action

- To clear a specified ACI interface set that that has no active members:

```
user@host> clear auto-configuration interfaces interface-set interface-set-name
```

For example, the following command clears the ACI interface set named aci-1003-ge-1/0/0.4001:

```
user@host> clear auto-configuration interfaces interface-set aci-1003-ge-1/0/0.4001
Interface-set aci-1003-ge-1/0/0.4001 deleted
```

Meaning The router dynamically creates an ACI interface set, if configured, when the first DHCP or PPPoE subscriber from a particular household logs in. However, the router does not automatically delete the ACI interface set when the last subscriber from that household logs out. As a result, you must use the **clear auto-configuration interfaces interface-set** command to explicitly clear the ACI interface set when it no longer has any active subscriber interface members. If you attempt to clear an ACI interface that still has active member interfaces, the router displays an error message and rejects the command.

When you specify the name of the ACI interface set to be cleared, you must use the ACI interface set name internally generated by the router, and not the actual ACI string carried in DHCP and PPPoE control packets. The router uses the following format to name ACI interface sets, as shown in the ACI interface set named aci-1003-ge-1/0/0.4001:

aci-*nnnn-interface-name.logical-unit-number*

 Copyright © 2013, Juniper Networks, Inc.

where:

- *nnnn* is a randomly generated 4-digit identifier (1003 in the example)
- *interface-name* is the name of the dynamic subscriber interface (ge-1/0/0 in the example)
- *logical-unit-number* is the logical unit number of the dynamic subscriber interface (4001 in the example)

To view the names of the ACI interface sets configured on the router, use the **show subscribers** command.

Related Documentation

- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration on page 747
- *Junos OS Operational Mode Commands*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 39

Dynamic VLAN Examples

- Example: Configuring a VLAN Dynamic Profile for VLANs with a TPID of 0x8100 on page 751
- Example: Configuring a VLAN Dynamic Profile for VLANs with Any TPID Value and Enabling Demux Interfaces over the VLAN Interface on page 751
- Example: Configuring a Stacked VLAN Dynamic Profile on page 752
- Example: Dynamic VLAN Interface Configuration on page 752
- Example: Dynamic Stacked VLAN Interface Configuration on page 753
- Example: Dynamic Flexible VLAN Interface Configuration on page 753
- Example: Configuring a Flexible VLAN Interface for Use with a Nonstandard Ethertype on page 754

Example: Configuring a VLAN Dynamic Profile for VLANs with a TPID of 0x8100

```
vlan-prof1 {
  interfaces {
    "$junos-interface-ifd-name" {
      unit "$junos-interface-unit" {
        vlan-id "$junos-vlan-id"; #Note the statement and variable use.
        family inet {
          unnumbered-address lo0.0 preferred-source-address 10.20.0.2;
        }
      }
    }
  }
}
```

Related Documentation

- Dynamic 802.1Q VLAN Overview on page 701
- Configuring a VLAN Dynamic Profile for Creating Single-Tag VLANs Using Standard TPID Values on page 714

Example: Configuring a VLAN Dynamic Profile for VLANs with Any TPID Value and Enabling Demux Interfaces over the VLAN Interface

```
vlan-prof-any-tpid {
  interfaces {
    $junos-interface-ifd-name {
```

```
            unit $junos-interface-unit {
              demux-source inet; #Enables demux interface use over the VLAN interface.
              vlan-tags outer $junos-vlan-id; #Statement/variable combination enables the
                recognition of any VLAN interface TPID value.
              family inet {
                unnumbered-address lo0.0 preferred-source-address 10.20.0.2;
              }
            }
          }
        }
      }
```

Related Documentation

- Dynamic 802.1Q VLAN Overview on page 701
- Configuring a VLAN Dynamic Profile for Creating Single-Tag VLANs Using Any TPID Values on page 716

Example: Configuring a Stacked VLAN Dynamic Profile

```
svlan-prof1 {
  interfaces {
    $junos-interface-ifd-name {
      unit $junos-interface-unit {
        vlan-tags outer $junos-stacked-vlan-id inner $junos-vlan-id;
        family inet {
          unnumbered-address lo0.0 preferred-source-address 100.20.0.2;
        }
      }
    }
  }
}
```

Related Documentation

- Dynamic 802.1Q VLAN Overview on page 701
- Configuring a Stacked VLAN Dynamic Profile on page 717

Example: Dynamic VLAN Interface Configuration

```
interfaces {
  ge-0/0/0 {
    vlan-tagging;
    auto-configure {
      vlan-ranges {
        dynamic-profile vlan-prof1 {
          accept inet;
          ranges {
            any;
          }
        }
      }
    }
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring VLAN Ranges for Use with Dynamic Profiles on page 724

Example: Dynamic Stacked VLAN Interface Configuration

```
interfaces {
  ge-0/0/0 {
    stacked-vlan-tagging;
    auto-configure {
      stacked-vlan-ranges {
        dynamic-profile svlan-prof {
          accept inet;
          ranges {
            1-1, any;
          }
        }
      }
    }
  }
}
```

Related Documentation

- Configuring Stacked VLAN Ranges for Use with Stacked VLAN Dynamic Profiles on page 725

Example: Dynamic Flexible VLAN Interface Configuration

```
interfaces {
  ge-0/0/0 {
    flexible-vlan-tagging;
    auto-configure {
      vlan-ranges {
        dynamic-profile vlan-prof1 {
          accept inet;
          ranges {
            any;
          }
        }
      }
      stacked-vlan-ranges {
        dynamic-profile svlan-prof1 {
          accept inet;
          ranges {
            1-1, any;
          }
        }
      }
    }
  }
}
```

Related Documentation

- Configuring Dynamic Mixed VLAN Ranges on page 726

Example: Configuring a Flexible VLAN Interface for Use with a Nonstandard Ethertype

This example specifies an ethertype of 0x9100 instead of the standard 0x8100.

```
interfaces {
  ge-0/0/0 {
    flexible-vlan-tagging;
    gigether-options {
      ethernet-switch-profile {
        tag-protocol-id 0x9100;
      }
    }
    auto-configure {
      vlan-ranges {
        dynamic-profile vlan-prof {
          accept inet;
          ranges {
            any;
          }
        }
      }
      stacked-vlan-ranges {
        dynamic-profile svlan-prof {
          accept inet;
          ranges {
            1-1,any;
          }
        }
      }
    }
  }
}
```

Related Documentation

- Dynamic 802.1Q VLAN Overview on page 701+
- *Configuring Inner and Outer TPIDs and VLAN IDs*

Copyright © 2013, Juniper Networks, Inc.

PART 12

Subscriber Interfaces

- Subscriber Interface Overview on page 757
- Configuring Subscriber Interfaces for Dynamic Profiles on page 765
- Subscriber Interface Examples on page 779
- Subscriber Interfaces over Aggregated Ethernet Overview on page 815
- Configuring Subscriber Interfaces over Aggregated Ethernet on page 823
- Subscriber Interfaces over Aggregated Ethernet Examples on page 835
- Dynamic PPPoE Subscriber Interfaces Overview on page 883
- Configuring Dynamic PPPoE Subscriber Interfaces on page 899
- Dynamic PPPoE Subscriber Interfaces Examples on page 917
- MPLS Pseudowire Subscriber Interfaces Overview on page 923
- Configuring MPLS Pseudowire Subscriber Interfaces on page 933

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 40

Subscriber Interface Overview

- Subscriber Interface Overview on page 757
- Static Subscriber Interfaces and VLAN Overview on page 758
- Subscriber Interfaces and Demultiplexing Overview on page 759
- IP Demux Interfaces over Static or Dynamic VLAN Demux Interfaces on page 760
- MAC Address Validation for Subscriber Interfaces Overview on page 761

Subscriber Interface Overview

In this release, you can identify subscribers statically or dynamically.

To identify subscribers statically, you can reference a static VLAN interface in a dynamic profile. To identify subscribers dynamically, you create variables for demux interfaces that are dynamically created by DHCP when subscribers log in.

Statically Identifying Subscribers

Before you can configure static subscriber interfaces in a dynamic profile, you must first configure the logical interfaces on the router to which you expect clients to connect. After you have created the static interfaces, you can modify them by using dynamic profiles to apply configuration parameters.

You can also configure subscribers by creating sets of static IP demux interfaces that are not referenced in a dynamic profile.

When configuring the interfaces stanza within a dynamic profile, you use variables to specify the interface name and the logical unit value. When a DHCP subscriber sends a DHCP request to the interface, the dynamic profile replaces the **interface-name** and **unit** variables with the actual interface name and logical unit number of the interface that received the DHCP request. After this association is made, the router configures the interface with any CoS or protocol (that is, IGMP) configuration within the dynamic profile, or applies any input or output filter configuration that you have associated with that dynamic profile.

```
[edit dynamic-profiles]
interfaces interface-name {
  unit logical-unit-number {
    family family {
      address address;
```

```
            filter {
                input filter-name;
                output filter-name;
            }
        unnumbered-address interface-name <preferred-source-address address>;
        vlan-id;
    }
    vlan-tagging;
}
```

Dynamically Identifying Subscribers

You can configure demux interfaces to represent a subscriber interface in a dynamic profile. When a subscriber logs in using a DHCP access method, the demux interface is dynamically created.

You specify variables for the unit number, the name of the underlying interface, and the IP address in the dynamic profile. These variables are replaced with the values that are supplied by DHCP when the subscriber logs in.

Related Documentation

- Static Subscriber Interfaces and VLAN Overview on page 758
- Subscriber Interfaces and Demultiplexing Overview on page 759

Static Subscriber Interfaces and VLAN Overview

This topic describes the topology for configuring subscriber interfaces over static VLAN interfaces in the current release.

In a dynamic profile, you can configure VLAN subscriber interfaces over the following statically created logical interface types:

- GE—Gigabit Ethernet
- XE—10-Gigabit Ethernet
- AE—Aggregated Ethernet

We recommend that you configure each subscriber on a statically created VLAN.

Figure 13 on page 758 shows an example of subscriber interfaces on an individual VLAN.

Figure 13: VLAN Subscriber Interfaces

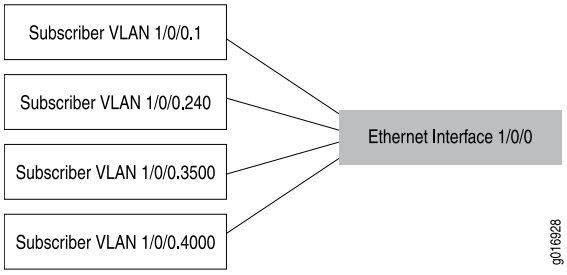

Copyright © 2013, Juniper Networks, Inc.

You can further separate VLANs on subscriber interfaces by configuring a VLAN interface as the underlying interface for a set of IP demux interfaces.

Related Documentation

- Configuring a Subscriber Interface with a Static VLAN Interface on page 766
- For more information about demux interfaces, see Subscriber Interfaces and Demultiplexing Overview on page 759

Subscriber Interfaces and Demultiplexing Overview

You can create logical subscriber interfaces using static or dynamic demultiplexing interfaces. In addition, you can use either IP demultiplexing interfaces or VLAN demultiplexing interfaces when creating logical subscriber interfaces.

Demultiplexing (demux) interfaces are logical interfaces that share a common, underlying logical interface (in the case of IP demux) or underlying physical interface (in the case of VLAN demux). You can use these interfaces to identify specific subscribers or to separate individual circuits by IP address (IP demux) or VLAN ID (VLAN demux).

The subscriber interfaces can provide different levels of services for individual subscribers in an access network. For example, you can apply CoS parameters for each subscriber.

Interface Sets of Static Demux Interfaces

You can group static demux interfaces to create individual subscriber interfaces using interface sets. Interface sets enable you to provide the same level of service for a group of subscribers; for example, all residential subscribers who receive the basic data service.

Figure 14 on page 759 shows a subscriber interface configured using a set of IP demux interfaces with an underlying VLAN interface.

Figure 14: IP Demux Subscriber Interface

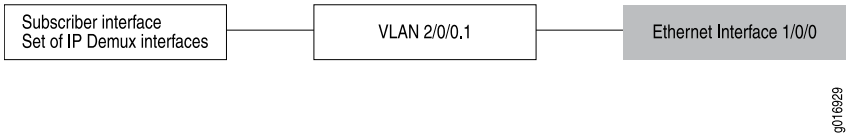

Dynamic Demultiplexing Interfaces

You can configure demux interfaces to represent a dynamic subscriber interface in a dynamic profile.

Demux interfaces are dynamically created by a DHCP access method when the underlying interface for the demux interface is configured for the access method. The DHCP access model creates the demux interface with the subscriber's assigned IP address (for IP demux interfaces) or VLAN ID (for VLAN demux interfaces).

To configure an IP demux interface in the dynamic profile, you specify variables for the unit number, the name of the underlying interface, and the IP address. To configure a VLAN demux interface in the dynamic profile, you specify variables for the unit number,

Copyright © 2013, Juniper Networks, Inc.

the name of the underlying interface, and the VLAN ID. These variables are replaced with the values that are supplied by DHCP when the subscriber logs in.

Guidelines for Configuring Demux Interfaces for Subscriber Access

When you configure static or dynamic demux interfaces for subscriber access, consider the following guidelines:

- You can only configure interface sets of static demux interfaces and dynamic demux interfaces on MX Series 3D Universal Edge Routers. Hierarchical and per-unit scheduling is supported for dynamically created demux interfaces on the EQ DPC.
- You can configure IPv4 and IPv6 addressing for static and dynamic demux interfaces.
- You can configure only one **demux0** interface per chassis.
- For IP demux interfaces, you can define logical demux interfaces on top of the **demux0** interface (for example, **demux0.1**, **demux0.2**, and so on).
- Demux interfaces currently support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet underlying interfaces.
- You must associate IP demux interfaces with an underlying logical interface.
- You must associate VLAN demux interfaces with an underlying device (physical interface).
- You cannot use a dynamic demux interface to represent multiple subscribers in a dynamic profile attached to an interface. One dynamic demux interface represents one subscriber. Do not configure the **aggregate-clients** option when attaching a dynamic profile to a demux interface for DHCP.

Related Documentation

- Configuring Static Subscriber Interfaces Using IP Demux Interfaces on page 766
- Configuring Static Subscriber Interfaces Using VLAN Demux Interfaces on page 767
- Configuring a Subscriber Interface Using a Set of Static IP Demux Interfaces on page 769
- Configuring a Subscriber Interface Using a Set of Static VLAN Demux Interfaces on page 770
- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles on page 772
- *CoS and Static IP Demux Interface Set Overview*
- For more information about static demux interfaces and other configuration guidelines, see the *Junos OS Network Interfaces Library for Routing Devices*

IP Demux Interfaces over Static or Dynamic VLAN Demux Interfaces

You can configure a router with IP demux interfaces over VLAN demux interfaces. Just as IP demux interfaces demultiplex their underlying VLAN demux interfaces based on IP

 Copyright © 2013, Juniper Networks, Inc.

address, VLAN demux interfaces demultiplex their underlying aggregate Ethernet or Ethernet interfaces based on VLAN ID.

When configuring IP demux interfaces over VLAN demux interfaces, keep the following in mind:

- Only single and dual VLAN tag options are supported as VLAN selectors.
- Both inet and inet6 families are supported.
- All firewall and CoS features are supported.
- Both static and dynamic VLAN demux interface creation is supported, including autosense VLAN creation.
- Both DPCs and MPCs are supported.

Related Documentation

- Subscriber Interfaces and Demultiplexing Overview on page 759
- Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 818
- Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet on page 825
- Configuring VLAN Dynamic Profiles on page 714
- Example: Dynamic IP Demux Subscriber Interfaces over Dynamic VLAN Demux Interfaces on page 795
- Example: Concurrent Configuration of Dynamic DHCP IP Demux and PPPoE Demux Interfaces over the Same VLAN Demux Interface on page 801
- For more information about aggregated Ethernet interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

MAC Address Validation for Subscriber Interfaces Overview

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address.

Configuring MAC address validation can provide additional validation when subscribers access billable services. MAC address validation provides additional security by enabling the router to drop packets that do not match, such as packets with spoofed addresses.

When subscribers log in, they are automatically assigned IP addresses by DHCP. With MAC address validation enabled, the router compares the IP source and MAC source addresses against trusted addresses, and forwards or drops the packets according to the match and the validation mode.

Copyright © 2013, Juniper Networks, Inc.

Supported Types of Subscriber Interfaces

MAC address validation is supported on statically or dynamically created Ethernet interfaces and demux interfaces on MX Series 3D Universal Edge Routers as follows:

- When the router is configured for a normal (non-enhanced) network services mode, MAC address validation is supported on both DPCs and MPCs. The router can be populated completely with one or the other type of line card, or have a mix of both types. Normal network services mode is the default.
- When the router is configured for Enhanced IP Network Services mode or Enhanced Ethernet Network Services mode, MAC address validation is supported only on MPCs. If the router has both DPCs and MPCs, or only DPCs, you cannot configure the chassis to be in enhanced mode.

MAC address validation is optimized for scaling when the router is in enhanced network services modes. Enhanced network services modes affect other features, such as multicast and firewall filters, so you must take that in to consideration when deciding whether to configure enhanced mode. For more information about the enhanced network service modes, see *Network Services Mode Overview* in the *Junos OS System Basics Configuration Guide*.

In normal network services mode, you can use the **show interfaces statistics** ***interface-name*** command to display a per-interface count of the packets that failed validation and were dropped. In enhanced network services mode, this command does not count the dropped packets; you must contact Juniper Networks Customer Support for assistance in collecting this data.

Trusted Addresses

A trusted address tuple is a 32–bit IP address and a 48–bit MAC address. Prefixes and ranges are not supported.

The IP source address and the MAC source address used for validation must be from a trusted source.

All static ARP addresses configured through the CLI are trusted addresses; dynamic ARP addresses are not considered trusted addresses.

Addresses dynamically created through an extended DHCP local server or extended DHCP relay are also trusted addresses. When a DHCP server and client negotiate an IP address, the resulting IP address and MAC address tuple is trusted. Each DHCP subscriber can generate more than one address tuple.

Each MAC address can have more than one IP address, which can result in more than one valid tuple. Each IP address must map to one MAC address.

Types of MAC Address Validation

You can configure either of two types or modes of MAC address validation, loose or strict. The behavior of the two modes varies depending on how well the incoming packets match the trusted address tuples. The modes differ only when the IP source address

 Copyright © 2013, Juniper Networks, Inc.

alone does not match any trusted IP address. Table 68 on page 763 compares the behavior of the two modes. Dropped packets are considered to be spoofed.

Table 68: Comparison of MAC Address Validation Modes

Incoming Packet Addresses Match Trusted Address Tuple	Loose Mode Action	Strict Mode Action
• IP source address matches and • MAC source address matches	Forwards packet	Forwards packet
• IP source address matches but • MAC source address does not match	Drops packet	Drops packet
• IP source address does not match and • MAC source address either matches or does not match	Forwards packet	Drops packet

Configuring strict mode is a more conservative strategy because it requires both received source addresses to match trusted addresses.

When you configure MAC address validation for IP demux interfaces in a dynamic profile and specify either loose or strict validation, the resulting behavior is always loose validation. To enable strict behavior for a dynamic IP demux interface, you must configure strict validation for both the IP demux interface and the underlying interface.

Related Documentation

- Configuring MAC Address Validation for Subscriber Interfaces on page 773

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 41

Configuring Subscriber Interfaces for Dynamic Profiles

- Configuring Static Subscriber Interfaces in Dynamic Profiles on page 765
- Configuring a Subscriber Interface Using a Set of Static IP Demux Interfaces on page 769
- Configuring a Subscriber Interface Using a Set of Static VLAN Demux Interfaces on page 770
- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles on page 772
- Configuring MAC Address Validation for Subscriber Interfaces on page 773
- Verifying Configuration and Status of Dynamic Subscribers and Associated Sessions, Services, and Firewall Filters on page 776

Configuring Static Subscriber Interfaces in Dynamic Profiles

In this release, you can use dynamic profiles to configure statically created logical interfaces. Dynamic profiles enable you to dynamically apply configured values (including CoS, IGMP, or filter configuration) to the static interfaces, making them easier to manage.

To configure static interfaces, you must first configure the interfaces on the router to which you expect subscribers to connect.

The subscriber access feature supports the following statically-created interface types in dynamic profiles:

- GE—Gigabit Ethernet
- XE—10-Gigabit Ethernet
- AE—Aggregated Ethernet

This topic contains the following sections:

- Configuring a Subscriber Interface with a Static VLAN Interface on page 766
- Configuring Static Subscriber Interfaces Using IP Demux Interfaces on page 766

Copyright © 2013, Juniper Networks, Inc.

- Configuring Static Subscriber Interfaces Using VLAN Demux Interfaces on page 767
- Associating Dynamic Profiles with Statically Created Interfaces on page 768

Configuring a Subscriber Interface with a Static VLAN Interface

This topic describes how to configure a subscriber interface with a static VLAN interface.

After you configure the static VLAN interface, you can reference it in a dynamic profile.

To configure a subscriber interface over a VLAN:

1. Configure the static VLAN interface and enable VLAN tagging.

```
[edit interfaces]
  ge-5/0/0 {
    vlan-tagging;
  }
```

2. Configure the units and assign the VLAN IDs.

```
unit 1 {
  proxy-arp;
  vlan-id 1;
  family inet {
    unnumbered-address lo0.0 preferred-source-address 192.1.1.1;
  }
}
unit 2 {
  proxy-arp;
  vlan-id 2;
  family inet {
    unnumbered-address lo0.0 preferred-source-address 192.1.1.1;
  }
}
```

3. Associate the static subscriber interface in a dynamic profile.

   See "Associating Dynamic Profiles with Statically Created Interfaces" on page 768.

Configuring Static Subscriber Interfaces Using IP Demux Interfaces

You can configure a subscriber interface using a statically created IP demux interface. This interface can be referenced in a dynamic profile.

To configure a static IP demux subscriber interface:

1. Configure the IP demux interface on a physical device represented by a logical unit number. The logical interface resides on a physical device.

   See *Configuring an IP Demultiplexing Interface*.

2. Configure the underlying interface on which the IP demux interface is running.

   See *Configuring an IP Demux Underlying Interface*

3. Specify the underlying interface on which the IP demux interface is running.

   See *Specifying the Demux Underlying Interface*.

 Copyright © 2013, Juniper Networks, Inc.

4. Specify how ingress IPv4 traffic is to be demultiplexed based on packet destination or source addresses.

   See *Configuring IP Demux Prefixes*.

5. Associate the static subscriber interface in a dynamic profile.

   See "Associating Dynamic Profiles with Statically Created Interfaces" on page 768.

NOTE: **VLAN demux interfaces currently support the Internet Protocol version 4 (IPv4) suite (family inet) and the Internet Protocol version 6 (IPv6) suite (family inet6).**

Configuring Static Subscriber Interfaces Using VLAN Demux Interfaces

You can configure a subscriber interface using a statically created VLAN demux interface. This interface can be referenced in a dynamic profile.

To configure a static VLAN demux subscriber interface:

1. Configure the VLAN demux interface.

   See *Configuring a VLAN Demultiplexing Interface*.

2. Configure the underlying interface on which the VLAN demux interface is running.

   See *Configuring a VLAN Demux Underlying Interface*

3. Specify the underlying interface on which the VLAN demux interface is running.

   See *Specifying the Demux Underlying Interface*.

4. Specify how ingress IP traffic is to be demultiplexed based on the VLAN ID.

   See *Associating VLAN IDs to VLAN Demux Interfaces*.

5. Associate the static subscriber interface in a dynamic profile.

   See "Associating Dynamic Profiles with Statically Created Interfaces" on page 768.

NOTE: **VLAN demux interfaces currently support the Internet Protocol version 4 (IPv4) suite (family inet) and the Internet Protocol version 6 (IPv6) suite (family inet6).**

VLAN demux subscriber interfaces over aggregated Ethernet physical interfaces are supported only for MX Series routers that have only MPCs installed. If the router has other cards in addition to MPCs, the CLI accepts the configuration but errors are reported when the subscriber interfaces are brought up.

Associating Dynamic Profiles with Statically Created Interfaces

When configuring the interfaces stanza within a dynamic profile, you use variables to specify the interface name and the logical unit value. When a DHCP subscriber sends a DHCP request to the interface, the dynamic profile replaces the interface name variable and logical unit name variable with the actual interface name and logical unit number of the interface that received the DHCP request.

NOTE: **Configuration of the interface name variable and logical interface name variable at the [edit dynamic-profiles *profile-name* interfaces] hierarchy level is required for a dynamic profile to function.**

To configure the interface for a dynamic profile, specify the interface name variable and include the **unit** statement and associated logical interface name variable:

1. Access the profile.

   ```
   [edit]
   user@host# edit dynamic-profiles basic-profile
   ```

2. Specify the interface name variable.

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

3. Specify the logical interface name variable with the **unit** statement.

   When referencing an existing interface, specify the **$junos-underlying-interface-unit** variable used by the router to match the unit value of the receiving interface:

   ```
   [edit dynamic-profiles basic-profile]
   user@host# set unit $junos-underlying-interface-unit
   ```

   When creating dynamic interfaces, specify the **$junos-interface-unit** variable used by the router to generate a unit value for the interface:

   ```
   [edit dynamic-profiles basic-profile]
   user@host# set unit $junos-interface-unit
   ```

Related Documentation

- Static Subscriber Interfaces and VLAN Overview on page 758
- For information about configuring logical interfaces and static VLAN interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

Configuring a Subscriber Interface Using a Set of Static IP Demux Interfaces

You can create logical subscriber interfaces from IP demux interfaces. IP demultiplexing (demux) interfaces are logical interfaces that share a common, underlying logical interface. IP demux interfaces can be used to identify specific subscribers or to separate individual circuits.

You can group individual subscriber interfaces using interface sets to provide the same level of service for a group of subscribers; for example, all residential subscribers who receive the basic data service. Interface sets can be defined as a list of logical interfaces (unit 0, unit 1, and so on).

To configure a group of static IP demux interfaces:

1. Configure the interface set.

```
interfaces {
  interface-set demux-set {
    interface demux0 {
      unit 0;
      unit 1;
      }
    }
  }
```

2. Define the units of the interface set.

```
demux0 {
  unit 0 {
    demux-options {
      underlying-interface ge-2/0/1.1;
    }
    family inet {
      demux-source {
        1.1.1.0/24;
      }
      address 1.1.1.1/24;
    }
  }
  unit 1 {
    demux-options {
      underlying-interface ge-2/0/1.1;
    }
    family inet {
      demux-source {
        1.1.2.0/24;
      }
      address 1.1.2.1/24;
    }
  }
}
```

Related Documentation

- *Configuring CoS on a Set of Static IP Demux Interfaces*
- Subscriber Interfaces and Demultiplexing Overview on page 759

- For information about the **[edit interfaces]** hierarchy and the **interface-set** statement, see the *Junos OS Network Interfaces Library for Routing Devices*

Configuring a Subscriber Interface Using a Set of Static VLAN Demux Interfaces

You can create logical subscriber interfaces from VLAN demux interfaces. VLAN demultiplexing (demux) interfaces are logical interfaces that share a common, underlying physical interface. VLAN demux interfaces can be used to identify specific subscribers or to separate individual circuits.

You can group individual subscriber interfaces using interface sets to provide the same level of service for a group of subscribers; for example, all residential subscribers who receive the basic data service. Interface sets can be defined as a list of logical interfaces (unit 0, unit 1, and so on).

To configure a group of static VLAN demux interfaces:

1. Configure the interface set.

```
interfaces {
  interface-set demux-set {
    interface demux0 {
      unit 0;
      unit 1;
      }
    }
  }
```

2. Define the units of the interface set.

```
demux0 {
  unit 0 {
    vlan-id 10;
    demux-options {
      underlying-interface ge-2/0/1;
    }
    family inet {
      address 1.1.1.1/24;
    }
  }
  unit 1 {
    vlan-id 20;
    demux-options {
      underlying-interface ge-2/0/1;
    }
    family inet {
      address 1.1.2.1/24;
    }
  }
}
```

Related Documentation

- *Configuring CoS on a Set of Static IP Demux Interfaces*
- Subscriber Interfaces and Demultiplexing Overview on page 759

 Copyright © 2013, Juniper Networks, Inc.

- For information about the **[edit interfaces]** hierarchy and the **interface-set** statement, see the *Junos OS Network Interfaces Library for Routing Devices*

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

You can configure dynamic subscriber interfaces using IP demux interfaces.

To enable the dynamic demux interface to be created by DHCP, you configure the demux options in a dynamic profile. Dynamic profiles enable you to dynamically apply configured values (including CoS, IGMP, or filter configuration) to the dynamic interfaces, making them easier to manage.

Before you begin:

- Configure the dynamic profile.

  See "Configuring a Basic Dynamic Profile" on page 673.

To configure dynamic subscriber interfaces:

1. Specify that you want to configure the **demux0** interface in the dynamic profile.

   ```
   user@host# edit dynamic-profiles business-profile interfaces demux0
   ```

2. Configure the unit for the **demux0** interface.

   a. Configure the variable for the unit number of the **demux0** interface.

      The variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

      ```
      [edit dynamic-profiles business-profile interfaces demux0]
      user@host# edit unit $junos-interface-unit
      ```

   b. Configure the variable for the underlying interface of the demux interfaces and specify the **$junos-underlying-interface** variable.

      The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.

      ```
      [edit dynamic-profiles business-profile interfaces demux0 unit
        "$junos-interface-unit"]
      user@host# set demux-options underlying-interface $junos-underlying-interface
      ```

3. Configure the family for the demux interfaces.

   a. Specify that you want to configure the family.

      For IPv4:

      ```
      [edit dynamic-profiles business-profile interfaces demux0 unit
        "$junos-interface-unit"]
      user@host# edit family inet
      ```

      For IPv6:

```
[edit dynamic-profiles business-profile interfaces demux0 unit
   "$junos-interface-unit"]
user@host# edit family inet6
```

b. Configure the unnumbered address for the family.

```
[edit dynamic-profiles business-profile interfaces demux0 unit
   "$junos-interface-unit" family inet]
user@host#  set unnumbered-address lo0.0
```

c. Configure the variable for the IP address of the demux interface.

The variable is dynamically replaced with the IP address that DHCP supplies when the subscriber logs in. For IPv4, use **$junos-subscriber-ip-address**, For IPv6, use **$junos-subscriber-ipv6-address**. For IPv6 multiple address support, use **$junos-subscriber-ipv6-multi-address**.

```
[edit dynamic-profiles business-profile interfaces demux0 unit
   "$junos-interface-unit" family inet]
user@host# set demux-source $junos-subscriber-ip-address
```

Related Documentation

- Subscriber Interfaces and Demultiplexing Overview on page 759
- Configuring MAC Address Validation for Dynamic Subscriber Interfaces on page 775
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236
- Example: Configuring Dynamic Subscriber Interfaces on IP Demux Interfaces on page 781

Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles

You can configure dynamic subscriber interfaces using VLAN demux interfaces.

To enable the dynamic demux interface to be created by DHCP, you configure the demux options in a dynamic profile. Dynamic profiles enable you to dynamically apply configured values (including CoS, IGMP, or filter configuration) to the dynamic interfaces, making them easier to manage.

Before you begin:

- Configure the dynamic profile.

  See "Configuring a Basic Dynamic Profile" on page 673.

To configure dynamic subscriber interfaces:

1. Specify that you want to configure the **demux0** interface in the dynamic profile.

```
user@host#  edit dynamic-profiles business-profile interfaces demux0
```

2. Configure the unit for the **demux0** interface.

   a. Configure the variable for the unit number of the **demux0** interface.

Copyright © 2013, Juniper Networks, Inc.

The variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles business-profile interfaces demux0]
user@host# edit unit $junos-interface-unit
```

b. Configure the variable for the underlying interface of the demux interfaces by specifying the **$junos-interface-ifd-name** variable.

The variable is dynamically replaced with the underlying device name that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles business-profile interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set demux-options underlying-interface $junos-interface-ifd-name
```

c. Configure the variable for the VLAN ID.

```
[edit dynamic-profiles business-profile interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set vlan-id $junos-vlan-id
```

3. Configure the family for the demux interfaces.

a. Specify that you want to configure the family.

For IPv4:

```
[edit dynamic-profiles business-profile interfaces demux0 unit
  "$junos-interface-unit"]
user@host# edit family inet
```

For IPv6:

```
[edit dynamic-profiles business-profile interfaces demux0 unit
  "$junos-interface-unit"]
user@host# edit family inet6
```

b. Configure the unnumbered address for the family.

```
[edit dynamic-profiles business-profile interfaces demux0 unit
  "$junos-interface-unit" family inet]
user@host#  set unnumbered-address lo0.0
```

Related Documentation

- Subscriber Interfaces and Demultiplexing Overview on page 759
- Configuring MAC Address Validation for Dynamic Subscriber Interfaces on page 775
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236
- Example: Dynamic IP Demux Subscriber Interfaces over Dynamic VLAN Demux Interfaces on page 795

Configuring MAC Address Validation for Subscriber Interfaces

This topic describes how to configure MAC address validation for subscriber interfaces in dynamic profiles on MX Series routers.

The subscriber interfaces can be statically created and associated with a dynamic profile (for example, VLAN interfaces) or dynamically created in the dynamic profile (such as demux interfaces).

By default, MAC address validation is disabled.

This topic contains the following sections:

- Configuring MAC Address Validation for Static Subscriber Interfaces on page 774
- Configuring MAC Address Validation for Dynamic Subscriber Interfaces on page 775

Configuring MAC Address Validation for Static Subscriber Interfaces

This topic describes how to configure MAC address validation for static subscriber interfaces in dynamic profiles on MX Series routers.

Before you begin:

- Configure the dynamic profile.

  See "Configuring a Basic Dynamic Profile" on page 673.

- (Optional) Configure an enhanced network services mode.

  See *Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers*.

To configure MAC address validation on static subscriber interfaces:

1. Configure the static VLAN interface.

   [edit interfaces]
   user@host# **set** *interface-name* **unit** *logical-unit-number* **family inet**

2. Configure the type of MAC address validation for the interface.

   - To configure loose validation:

     [edit interfaces *interface-name* unit *logical-unit-number* family inet]
     user@host# **set mac-validate loose**

   - To configure strict validation:

     [edit interfaces *interface-name* unit *logical-unit-number* family inet]
     user@host# **set mac-validate strict**

For example, to configure loose validation on interface fe-0/0/0.0, configure the following:

[edit interfaces fe-0/0/0 unit 0 family inet]
user@host# **set mac-validate loose**

Copyright © 2013, Juniper Networks, Inc.

After you configure MAC address validation:

- Associate the static VLAN interface with the dynamic profile.

  See "Associating Dynamic Profiles with Statically Created Interfaces" on page 768.

Configuring MAC Address Validation for Dynamic Subscriber Interfaces

This topic describes how to configure MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

When you configure MAC address validation for demux interfaces in a dynamic profile and specify either **loose** or **strict** validation, the resulting behavior is always loose validation. To enable strict behavior for a dynamic IP demux interface, besides configuring either **loose** or **strict** mode on the IP demux interface, you must also configure strict validation on the underlying interface.

Before you begin:

- Configure the dynamic profile.

  See "Configuring a Basic Dynamic Profile" on page 673.

- Configure the dynamic IP demux interface.

  See "Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles" on page 771.

- (Optional) Configure an enhanced network services mode.

  See *Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers*.

To configure loose MAC address validation for a dynamic subscriber interface:

- Configure loose validation for the demux interface.

  ```
  [edit dynamic-profiles profile-name interfaces demux0 unit "$junos-interface-unit"
    family inet]
  user@host# set mac-validate loose
  ```

  For loose validation, you do not need to configure MAC address validation on the underlying interface.

To configure strict MAC address validation for a dynamic subscriber interface:

1. Configure validation for the demux interface.

   ```
   [edit dynamic-profiles profile-name interfaces demux0 unit "$junos-interface-unit"
     family inet]
   user@host# set mac-validate validation-mode
   ```

   NOTE: Remember, although you must configure validation on the IP demux interface, it does not matter which mode you specify because the behavior is always loose.

2. Configure strict validation for the underlying interface.

   [edit interfaces *interface-name* unit *logical-unit-number* family inet]
   user@host# **set mac-validate strict**

   The underlying interface in this case is statically configured—for example, ge-1/0/0.1—and assigned to a DHCP configuration group that is associated with the dynamic profile. In a more complicated configuration, the underlying interface itself can be configured by a dynamic profile; in that case the validation is configured in the profile that creates the underlying interface.

Related Documentation

- MAC Address Validation for Subscriber Interfaces Overview on page 761
- Example: Configuring Dynamic Subscriber Interfaces on IP Demux Interfaces on page 781

Verifying Configuration and Status of Dynamic Subscribers and Associated Sessions, Services, and Firewall Filters

Purpose Verify configuration and status of dynamic subscribers, sessions, services, and firewall filters.

You can display information about subscribers in different ways, depending on the options you use with the **show subscriber** command. You can use details from one set of output with another command to display more detailed information of interest.

Action

- To display basic information for all subscribers:

```
user@host> show subscribers
Interface IP Address/VLAN ID User Name LS:RI
demux0.1073741824 0x8100.1500 0x8100.2900 test1@test.com default:testnet
demux0.1073741825 0x8100.1500 0x8100.2901 test1@test.com default:testnet
demux0.1073741826 0x8100.1500 0x8100.2902 test1@test.com default:testnet
demux0.1073741827 0x8100.1500 0x8100.2903 test1@test.com default:testnet
demux0.1073741826 172.16.200.6 test1@test.com default:testnet
demux0.1073741827 172.16.200.7 test1@test.com default:testnet
demux0.1073741824 172.16.200.8 test1@test.com default:testnet
demux0.1073741825 172.16.200.9 test1@test.com default:testnet
demux0.1073741828 0x8100.1500 0x8100.2910 test1@test.com default:default
demux0.1073741828 20.20.0.2 test1@test.com default:default
```

- To display more detailed information about a particular subscriber interface:

```
user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: test1@test.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:6e:56:01:04
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
```

 Copyright © 2013, Juniper Networks, Inc.

```
Type: DHCP
User Name: test1@test.com
IP Address: 172.16.200.6
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:6e:56:01:04
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
```

- To display traffic information for firewall filters.

```
user@host> show firewall
...
Filter: content-cb-in-demux0.1073741826-in
Counters:
Name     Bytes  Packets
__junos-dyn-service-counter  84336      1004

Filter: content-cb-out-demux0.1073741826-out
Counters:
Name     Bytes   Packets
__junos-dyn-service-counter      0       0
...
```

Instead of issuing successive commands to track the details for one subscriber interface, you can choose to display detailed information for all subscribers. However, the more subscribers you have, the more tedious it becomes to look through all the results for particular items of interest.

- To display detailed information for all subscribers:

  user@host> **show subscribers detail**

  user@host> **show subscribers extensive**

Meaning The output examples in this section show increasingly detailed information about dynamically created subscriber interfaces, including how many there are, what they are, and their characteristics; how many service sessions are active and what they are; whether firewall filters are attached to the sessions and what those filters are; and how much, if any, traffic is being filtered.

In the sample output shown here, the **show subscriber** command lists all the subscriber logical interfaces, including demux0.1073741826. You then display details about that interface and its associated subscribers with the **show subscribers interface demux0.1073741826 extensive** command. The Service Session Name fields for service sessions 25 and 26 in that output show two services are active on the interface, SUB-QOS and service-cb-content. The IPv4 Input Filter Name and the IPv4 Output Filter Name fields show that two filters have been applied to the service-cb-content session: content-cb-in-demux0.1073741826-in and content-cb-out-demux0.1073741826-out. You then use the **show firewalls** command to list the filters and see how much, if any, traffic is being filtered.

Related Documentation

- *Junos OS Operational Mode Commands*
- *Junos OS Operational Mode Commands*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 42

Subscriber Interface Examples

- Example: Configuring a Static Subscriber Interface on a Gigabit Ethernet VLAN Interface (Multiple Logical Units) on page 779
- Example: Configuring a Static Subscriber Interface on a Gigabit Ethernet VLAN Interface on page 780
- Example: Configuring a Static Subscriber Interface on a Gigabit Ethernet VLAN Interface (No Autonegotiation) on page 780
- Example: Configuring a Static Subscriber Interface with a Loopback on page 780
- Example: Configuring Dynamic Subscriber Interfaces on IP Demux Interfaces on page 781
- Example: Configuring IPv6 Addressing for a Dynamic IP Demux Interface over Static VLANs on page 783
- Example: Configuring IPv6 Addressing for a Dynamic IP Demux Interface over Dynamic VLANs on page 784
- Example: Configuring a Dynamic IP Demux Interface with Dual Stacking on page 787
- Example: Configuring IPv4 Static VLAN Demux Interfaces over a Gigabit Ethernet Underlying Interface with DHCP Local Server on page 791
- Example: Configuring IPv4 Dynamic VLAN Demux Interfaces over a Gigabit Ethernet Underlying Interface with DHCP Local Server on page 793
- Example: Dynamic IP Demux Subscriber Interfaces over Dynamic VLAN Demux Interfaces on page 795
- Example: Concurrent Configuration of Dynamic DHCP IP Demux and PPPoE Demux Interfaces over the Same VLAN Demux Interface on page 801
- Example: Configuring CoS on Static LSQ MLPPP Bundle Interfaces on page 811

Example: Configuring a Static Subscriber Interface on a Gigabit Ethernet VLAN Interface (Multiple Logical Units)

```
[edit interfaces]
ge-5/0/0 {
   vlan-tagging;
   unit 1 {
      proxy-arp;
      vlan-id 1;
      family inet {
         unnumbered-address lo0.0 preferred-source-address 192.1.1.1;
```

```
        }
      }
      unit 2 {
        proxy-arp;
        vlan-id 2;
        family inet {
          unnumbered-address lo0.0 preferred-source-address 192.1.1.1;
        }
      }
    }
```

Related Documentation

- Configuring Static Subscriber Interfaces in Dynamic Profiles on page 765

Example: Configuring a Static Subscriber Interface on a Gigabit Ethernet VLAN Interface

```
[edit interfaces]
ge-5/2/0 {
  vlan-tagging;
  unit 1 {
    vlan-id 1;
    family inet {
      address 192.2.1.1/24;
    }
  }
}
```

Related Documentation

- Configuring Static Subscriber Interfaces in Dynamic Profiles on page 765

Example: Configuring a Static Subscriber Interface on a Gigabit Ethernet VLAN Interface (No Autonegotiation)

```
[edit interfaces]
ge-5/1/9 {
  vlan-tagging;
  gigether-options {
    no-auto-negotiation;
  }
  unit 2004 {
    vlan-id 2004;
    family inet {
      address 222.0.0.1/22;
    }
  }
}
```

Related Documentation

- Configuring Static Subscriber Interfaces in Dynamic Profiles on page 765

Example: Configuring a Static Subscriber Interface with a Loopback

```
lo0 {
  unit 0 {
```


Copyright © 2013, Juniper Networks, Inc.

```
            family inet {
              address 192.1.1.1/32;
            }
          }
        }
```

Related Documentation

- Configuring Static Subscriber Interfaces in Dynamic Profiles on page 765

Example: Configuring Dynamic Subscriber Interfaces on IP Demux Interfaces

This example shows how to configure dynamic subscriber interfaces on IP demux interfaces. DHCP dynamically creates the demux interface when a subscriber logs in.

To configure subscribers on dynamic IP demux interfaces:

1. Configure the static VLAN as the underlying interface.

```
interfaces {
  ge-0/3/0 {
    vlan-tagging;
    unit 0 {
      vlan-id 0;
      demux-source inet;
      family inet {
        unnumbered-address lo0.0;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 90.1.1.1/24;
      }
    }
  }
}
```

2. Configure the creation of demux interfaces in the dynamic profile.

```
dynamic-profiles {
  subscriber-profile {
    interfaces {
      demux0 {
        "$junos-interface-ifd-name" {
          unit "$junos-interface-unit" {
            demux-options {
              underlying-interface "$junos-underlying-interface";
            }
            family inet {
              demux-source {
                $junos-subscriber-ip-address;
              }
              filter {
                input ingressFilter;
                output egressFilter;
```

```
                    }
                    mac-validate loose;
                }
            }
        }
    }
  }
}
```

3. Configure the access method to dynamically create the demux interface.

   DHCP relay is the access method used in this example.

```
forwarding-options {
  dhcp-relay {
    traceoptions {
      flag all;
    }
    server-group {
      router {
        100.20.42.1;
      }
      dynamic-profile subscriber-profile;
      active-server-group erx;
      group one {
        interface ge-0/0/2.0 upto ge-0/0/2.4000;
        interface-client-limit 200
      }
    }
  }
```

4. Configure the interface for DHCP.

```
interfaces {
  traceoptions {
    flag all;
  }
  ge-0/0/2 {
    unit 0 {
      demux-source inet;
      family inet {
        unnumbered-address lo0.0;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 100.20.32.2/32;
      }
    }
  }
}
```

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771

Copyright © 2013, Juniper Networks, Inc.

- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

Example: Configuring IPv6 Addressing for a Dynamic IP Demux Interface over Static VLANs

In this example, the network administrator configures IPv6 addressing for a dynamic IP demux interface over a group of static VLANs.

```
[edit]
dynamic-profiles {
  dhcp-demux-profile {
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          demux-options {
            underlying-interface "$junos-underlying-interface";
          }
          family inet6 {
            address 2001::1/64;
            demux-source {
              $junos-subscriber-ipv6-address;
            }
          }
        }
      }
    }
  }
}
system {
  services {
    dhcp-local-server {
      dhcpv6 {
        dynamic-profile dhcp-demux-prof;
        group vlan {
          interface ge-1/0/0.100;
        }
      }
    }
  }
}
interfaces {
  ge-1/0/0 {
    vlan-tagging;
    unit 100 {
      demux-source inet6;
      vlan-id 100;
      family inet6 {
        address 2001::1/64;
      }
    }
  }
}
```

```
access {
  address-assignment {
    pool dhcp {
      family inet6 {
        prefix 2001:0000:0000:0000::/64;
        range limits prefix-length 74;
      }
    }
  }
}
```

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771

Example: Configuring IPv6 Addressing for a Dynamic IP Demux Interface over Dynamic VLANs

In this example, the network administrator configures IPv6 addressing for a dynamic IP demux interface over a group of dynamic VLANs.

```
dynamic-profiles {
  vlan-profile {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
          vlan-id "$junos-vlan-id";
          demux-source inet6;
          family inet6 {
            unnumbered-address lo0.0 preferred-source-address ::100.20.32.2;
          }
        }
      }
    }
  }
  svlan-profile {
    interfaces {
      $junos-interface-ifd-name {
        unit $junos-interface-unit {
          demux-source inet6;
          vlan-tags outer $junos-stacked-vlan-id inner $junos-vlan-id;
          family inet6 {
            unnumbered-address lo0.0 preferred-source-address ::100.20.32.2;
          }
        }
      }
    }
  }
  dhcp-demux-prof {
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          demux-options {
            underlying-interface "$junos-underlying-interface";
          }
```


Copyright © 2013, Juniper Networks, Inc.

```
            family inet6 {
              demux-source {
                $junos-subscriber-ipv6-address;
              }
              unnumbered-address lo0.0 preferred-source-address 2001:db8:ffff:1::1;
            }
          }
        }
      }
    }
    all-profile {
      interfaces {
        $junos-interface-ifd-name {
          unit $junos-underlying-interface-unit {
          }
        }
      }
    }
    }
    }
    services {
      dhcp-local-server {
        traceoptions {
          file dhcp size 1g;
          flag all;
        }
          dhcpv6 {
            authentication {
              password delpref;
              username-include {
                user-prefix localpool;
              }
            }
            group one {
              authentication {
                password delpref;
                username-include {
                  user-prefix localpool;
                }
              }
              dynamic-profile dhcp-demux-prof use-primary all-profile;
              interface ge-0/0/3.0;
            }
          }
          group v6 {
            authentication {
              password delpref;
              username-include {
                user-prefix localpool;
              }
            }
            dynamic-profile dynamic-profile dhcp-demux-prof use-primary all-profile;
            interface ge-1/2/0.0;
          }
        }
      }
```

```
}
interfaces {
  ge-1/0/0 {
    vlan-tagging;
    auto-configure {
      vlan-ranges {
        dynamic-profile vlan-profile {
          accept inet6;
          ranges {
            any;
          }
        }
      }
    }
  }
  ge-1/2/0 {
    flexible-vlan-tagging;
    auto-configure {
      vlan-ranges {
        dynamic-profile vlan-profile {
          accept inet6;
          ranges {
            any;
          }
        }
      }
      stacked-vlan-ranges {
        dynamic-profile svlan-profile {
          accept inet6;
          ranges {
            any,any;
          }
        }
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 100.20.32.2/32;
      }
      family inet6 {
        address ::100.20.32.2/128;
      }
    }
  }
}
access {
  address-assignment {
    pool v6 {
      family inet6 {
        network 100.20.0.0/16;
        range limited {
          low 100.20.0.10;
          high 100.20.128.250;
        }
```


Copyright © 2013, Juniper Networks, Inc.

```
                dhcp-attributes {
                  maximum-lease-time 84600;
                }
              }
            }
          }
        }
```

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771

Example: Configuring a Dynamic IP Demux Interface with Dual Stacking

In this example, the network administrator configures IPv4 and IPv6 addressing for a dynamic IP demux interface with a group of underlying static VLANs.

```
[edit]
dynamic-profiles {
  dhcp-demux-prof {
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          demux-options {
            underlying-interface "$junos-underlying-interface";
          }
          family inet {
            demux-source {
              $junos-subscriber-ip-address;
            }
            unnumbered-address lo0.0 preferred-source-address 3.1.1.1;
          }
          family inet6 {
            demux-source {
              $junos-subscriber-ipv6-address;
            }
            unnumbered-address lo0.0 preferred-source-address 2001:db8:ffff:1::1;
          }
        }
      }
    }
  }
  all-profile {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit";
      }
    }
  }
}
services {
  dhcp-local-server {
    traceoptions {
      file dhcp size 1g;
      flag all;
```

```
    }
    dhcpv6 {
      authentication {
        password delpref;
        username-include {
          user-prefix localpool;
        }
      }
      group groupv6 {
        authentication {
          password delpref;
          username-include {
            user-prefix localpool;
          }
        }
        dynamic-profile dhcp-demux-prof use-primary all-profile;
        interface ge-0/0/3.0;
      }
    }
    group groupv4 {
      authentication {
        password delprefv4;
        username-include {
          user-prefix localpoolv4;
        }
      }
      dynamic-profile dhcp-demux-prof;
      interface ge-0/0/2.0;
    }
  }
  processes {
    general-authentication-service {
      traceoptions {
        file auth;
        flag all;
      }
    }
  }
  interfaces {
    ge-0/0/0 {
      unit 0 {
        proxy-arp;
        family inet6 {
          address 4ffe::1:1/48;
        }
      }
    }
    ge-0/0/1 {
      vlan-tagging;
      gigether-options {
        no-auto-negotiation;
      }
      unit 10 {
        vlan-id 10;
        family inet {
          address 100.10.0.2/24;
```

 Copyright © 2013, Juniper Networks, Inc.

```
            }
        }
    }
    ge-0/0/2 {
        unit 0 {
            demux-source inet;
            proxy-arp;
            family inet {
                unnumbered-address lo0.0 preferred-source-address 3.1.1.1;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            demux-source inet6;
            proxy-arp;
            family inet6 {
                unnumbered-address lo0.0 preferred-source-address 2001:db8:ffff:1::1;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 3.1.1.1/32;
            }
            family inet6 {
                address 2001:db8:ffff:1::1/128;
            }
        }
    }
}
access {
    radius-server {
        100.10.0.1 {
            port 1812;
            secret "$9$xs5-dsgoGDjqgo"; ## SECRET-DATA
        }
    }
    profile wash-test {
        accounting-order radius;
        authentication-order radius;
        radius {
            authentication-server 100.10.0.1;
            accounting-server 100.10.0.1;
        }
        accounting {
            order radius;
            accounting-stop-on-failure;
            accounting-stop-on-access-deny;
            update-interval 10;
            statistics time;
        }
    }
    address-assignment {
        pool v4ville {
```

```
                family inet {
                    network 3.1.1.0/24;
                    range testv4 {
                        low 3.1.1.3;
                        high 3.1.1.50;
                    }
                }
            }
            pool v6ville {
                family inet6 {
                    prefix 2001:db8:ffff::/48;
                    range test {
                        low 2001:db8:ffff:1::2/128;
                        high 2001:db8:ffff:1::ffff/128;
                    }
                }
            }
        }
    }
}
[edit]
dynamic-profiles {
    dhcp-demux-profile {
        interfaces {
            demux0 {
                unit "$junos-interface-unit" {
                    demux-options {
                        underlying-interface "$junos-underlying-interface";
                    }
                    family inet {
                        demux-source {
                            $junos-subscriber-ip-address;
                        }
                        unnumbered-address ge-0/0/0.0 preferred-source-address 1.1.1.2;
                    }
                    family inet6 {
                        demux-source {
                            $junos-subscriber-ipv6-address;
                        }
                        unnumbered-address ge-0/0/3.0 preferred-source-address ::22.22.22.2;
                    }
                }
            }
        }
    }
}
```

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring IPv4 Static VLAN Demux Interfaces over a Gigabit Ethernet Underlying Interface with DHCP Local Server

This example shows how to configure a static IPv4 VLAN demux interface with gigabit Ethernet as the underlying interface. DHCP Local Server configuration enables the association of subscribers to the VLAN demux interface by listing the gigabit Ethernet interface in the DHCP local server configuration.

To configure dynamic subscribers on VLAN demux interfaces:

1. Enable VLAN tagging on the underlying interface that you plan to use for the VLAN demux interfaces.

```
interfaces {
  ge-5/0/0 {
    vlan-tagging;
  }
}
```

2. Define the loopback interface.

```
interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 192.16.1.1/32;
      }
    }
  }
}
```

3. Define the demux interface.

```
interfaces {
  demux0 {
    unit 102 {
      proxy-arp;
      vlan-id 103;
      demux-options {
        underlying-interface ge-5/0/0;
      }
      family inet {
        unnumbered-address lo0.0 preferred-source-address 173.16.1.1;
      }
    }
  }
}
```

4. Configure a dynamic profile for subscriber access.

```
dynamic-profiles {
  user-profile {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit" {
          family inet;
```

```
            }
          }
        }
      }
    }
```

5. Configure the access method used to dynamically create the subscriber interfaces.

   The following stanza specifies the gigabit Ethernet interface (**ge-5/0/0.0**) for use with the dynamically created subscriber interfaces.

```
system {
  services {
    dhcp-local-server {
      group myDhcpGroup {
        authentication {
          password test;
          username-include {
            user-prefix igmp-user1;
          }
        }
        dynamic-profile user-profile;
        interface ge-5/0/0.0;
      }
    }
  }
}
```

   Instead of using the gigabit Ethernet interface, you can alternatively specify the specific demux interface (**demux0.102**) as the device to use with the subscriber interfaces as follows:

```
system {
  services {
    dhcp-local-server {
      group myDhcpGroup {
        authentication {
          password test;
          username-include {
            user-prefix igmp-user1;
          }
        }
        dynamic-profile user-profile;
        interface demux0.102;
      }
    }
  }
}
```

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring IPv4 Dynamic VLAN Demux Interfaces over a Gigabit Ethernet Underlying Interface with DHCP Local Server

This example shows how to configure the dynamic creation of IPv4 VLAN demux interfaces with gigabit Ethernet as the underlying interface. DHCP Local Server configuration enables the association of subscribers to the VLAN demux interface by listing the aggregated Ethernet interface in the DHCP local server configuration.

To configure dynamic subscribers on dynamic VLAN demux interfaces:

1. Enable VLAN tagging and VLAN auto-configuration on the underlying gigabit Ethernet interface that you plan to use for dynamically created VLAN demux interfaces.

```
interfaces {
  ge-5/0/0 {
    hierarchical-scheduler;
    vlan-tagging;
    auto-configure {
      vlan-ranges {
        dynamic-profile auto-vlanDemux-profile {
          accept inet;
          ranges {
            103-103;
          }
        }
      }
    }
  }
}
```

2. Define the loopback interface.

```
interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 192.16.1.1/32;
      }
    }
  }
}
```

3. Configure a dynamic profile for subscriber access.

```
dynamic-profiles {
  user-profile {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit" {
          family inet;
        }
      }
    }
  }
}
```

4. Configure a dynamic profile for VLAN demux interface creation.

```
dynamic-profiles {
  auto-vlanDemux-profile {
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          vlan-id "$junos-vlan-id";
          demux-options {
            underlying-interface "$junos-interface-ifd-name";
          }
          family inet {
            filter {
              input rate_limit;
              output rate_limit;
            }
            unnumbered-address lo0.0 preferred-source-address 192.16.1.1;
          }
        }
      }
    }
  }
}
```

5. Configure the access method used to dynamically create the subscriber interfaces. The following stanza specifies the gigabit Ethernet interface (**ge-5/0/0.0**) for use with the dynamically created subscriber interfaces.

```
system {
  services {
    dhcp-local-server {
      group myDhcpGroup {
        authentication {
          password test;
          username-include {
            user-prefix igmp-user1;
          }
        }
        dynamic-profile user-profile;
        interface ge-5/0/0.0;
      }
    }
  }
}
```

Instead of using the gigabit Ethernet interface, you can alternatively specify **demux0** as the device to use with the subscriber interfaces as follows:

NOTE: Because the demux interfaces and unit numbers are created dynamically, the unit number is not specified for the demux0 interface.

```
system {
  services {
    dhcp-local-server {
      group myDhcpGroup {
```


Copyright © 2013, Juniper Networks, Inc.

```
                authentication {
                    password test;
                    username-include {
                        user-prefix igmp-user1;
                    }
                }
                dynamic-profile user-profile;
                interface demux0;
            }
        }
    }
}
```

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

Example: Dynamic IP Demux Subscriber Interfaces over Dynamic VLAN Demux Interfaces

This example describes how to configure dynamic IP demux interfaces over dynamic VLAN demux interfaces. You can also configure dynamic IP demux interfaces over static VLAN interfaces. For information on how to configure static VLAN interfaces, see the *Ethernet Interfaces*.

- Requirements on page 795
- Overview on page 795
- Configuration on page 795
- Verification on page 800

Requirements

Before you begin, make sure to configure either DHCP Relay or DHCP Local Server. You can find information about configuring either of these components in the *Junos OS Subscriber Management and Services Library*.

Overview

You can create a subscriber interface using an IP demux interface stacked on a static or dynamic VLAN demux interface. IP demux interfaces are used to uniquely identify subscribers in an access network based on their IP address

Configuration

- Preparing a Subscriber Access Interface on page 796
- Preparing the Loopback Interface on page 797
- Configuring a Dynamic Profile to Dynamically Create Single-Tagged VLANs on page 798
- Configuring a Dynamic Profile to Dynamically Create IP Demux Interfaces on page 799

Preparing a Subscriber Access Interface

CLI Quick Configuration

To quickly configure the aggregated Ethernet interface over which subscribers access the router:

```
[edit]
set chassis aggregated-devices ethernet device-count 1
set interfaces ge-5/0/9 gigether-options 802.3ad ae0
set interfaces ge-5/1/9 gigether-options 802.3ad ae0
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 auto-configure vlan-ranges dynamic-profile Auto-VLAN-Demux accept
   inet
set interfaces ae0 auto-configure vlan-ranges dynamic-profile Auto-VLAN-Demux ranges
   ranges 500-1000
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp link-protection
```

Step-by-Step Procedure

You must configure an interface over which clients initially access the router. We recommend that you specify the same VLAN tagging for the interface that you expect from incoming clients. This example uses flexible VLAN tagging to simultaneously support transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port.

If you want it to automatically create dynamic VLANs, the interface must include the VLAN range type (single or stacked) and contain any specific ranges you want the VLANs to use.

To configure an interface for subscriber access:

1. Configure the number of aggregated Ethernet interfaces on the router.

   ```
   [edit]
   user@host# set chassis aggregated-devices ethernet device-count 1
   ```

2. Access the physical interface over which you want subscribers to initially access the router.

   ```
   [edit]
   user@host# edit interfaces ge-5/0/9
   ```

3. Specify the aggregated Ethernet interface to which the physical interface belongs.

   ```
   [edit interfaces ge-5/0/9]
   user@host# set gigether-options 802.3ad ae0
   ```

4. Repeat Step 2 and Step 3 for each interface you want to assign to the aggregated Ethernet bundle.

   ```
   [edit]
   user@host# set interfaces ge-5/1/9 gigether-options 802.3ad ae0
   ```

5. Access the aggregated Ethernet interface.

   ```
   [edit]
   user@host# edit interfaces ae0
   ```

6. Specify the VLAN tagging that you want the aggregated Ethernet interfaces to use.

Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces ae0]
user@host# set vlan-tagging
```

7. Edit the **auto-configure** stanza to automatically configure VLANs.

```
[edit interfaces ae0]
user@host# edit auto-configure
```

8. Edit the **vlan-ranges** stanza for single-tagged VLANs.

```
[edit interfaces ae0 auto-configure]
user@host# edit vlan-ranges
```

9. Specify the dynamic VLAN profile that you want the interface to use for dynamically creating single-tagged VLANs.

```
[edit interfaces ae0 auto-configure vlan-ranges]
user@host# edit dynamic-profile Auto-VLAN-Demux
```

10. Specify what VLAN Ethernet packet type the VLAN profile accepts.

```
[edit interfaces ae0 auto-configure vlan-ranges dynamic-profile Auto-VLAN-Demux]
user@host# set accept inet
```

11. Specify the VLAN ranges that you want the dynamic profile to use. The following example specifies a lower VLAN ID limit of 500 and an upper VLAN ID limit of 1000.

```
[edit interfaces ae0 auto-configure vlan-ranges dynamic-profile Auto-VLAN-Demux]
user@host# set ranges 500-1000
```

12. (Optional) Activate the transmission of LACP packets on the aggregated Ethernet interfaces.

```
[edit interfaces ae0]
user@host# set aggregated-ether-options lacp active
```

13. Specify that the aggregated Ethernet interfaces use link protection.

```
[edit interfaces ae0]
user@host# set aggregated-ether-options lacp link-protection
```

Preparing the Loopback Interface

CLI Quick Configuration

To quickly configure the required loopback interface for this example:

```
[edit]
set interfaces lo0.0 unit 0 family inet address 100.100.100.1/32
```

Step-by-Step Procedure

You must configure a loopback interface for use as the unnumbered address and preferred source address for dynamically created interfaces.

To configure the required loopback interface for this example:

1. Configure a loopback interface.

```
[edit]
user@host# edit interfaces lo0.0
```

2. Specify that the loopback interface accept inet packets.

```
[edit interfaces lo0 unit 0]
```

```
user@host# edit family inet
```

3. Specify the IP address for the loopback interface.

```
[edit interfaces lo0 unit 0 family inet]
user@host# set address 100.100.100.1/32
```

Configuring a Dynamic Profile to Dynamically Create Single-Tagged VLANs

CLI Quick Configuration

To quickly configure the dynamic profile used to dynamically create single-tagged VLANs in the example:

```
[edit]
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
  demux-source inet
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
  proxy-arp
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
  vlan-id $junos-vlan-id
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
  demux options underlying-interface $junos-interface-ifd-name
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
  family inet unnumbered-address lo0.0 preferred source-address 100.100.100.1
```

Step-by-Step Procedure

For dynamic IP demux interfaces to reside on a dynamic VLAN demux interface, the VLAN interface must first exist.

To configure a dynamic profile that automatically creates VLAN interfaces:

1. Create a dynamic profile for automatically creating single-tagged VLAN interfaces.

```
[edit]
user@host# edit dynamic-profiles Auto-VLAN-Demux
```

2. Specify that the dynamic VLAN profile use the demux interface.

```
[edit dynamic-profiles “Auto-VLAN-Demux”]
user@host# edit interfaces demux0
```

3. Specify that the dynamic profile apply the demux interface unit value to the dynamic VLANs.

```
[edit dynamic-profiles Auto-VLAN-Demux interfaces demux0]
user@host# edit unit $junos-interface-unit
```

4. (Optional) Specify that the demux source accepts only IPv4 (inet) packets.

```
[edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
  “$junos-interface-unit”]
user@host# set demux-source inet
```

5. (Optional) Specify that each dynamically created interface respond to any ARP request, as long as an active route exists to the target address of the ARP request.

```
[edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
  “$junos-interface-unit”]
user@host# set proxy-arp
```

6. Specify that VLAN IDs are dynamically created.

Copyright © 2013, Juniper Networks, Inc.

```
[edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set vlan-id $junos-vlan-id
```

7. Specify the logical underlying interface for the dynamic VLANs.

```
[edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set demux-options underlying-interface $junos-interface-ifd-name
```

8. Specify that the VLAN demux interface can accept inet family packets for IPoE/DHCP subscribers.

```
[edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# edit family inet
```

9. Specify the loopback address as the unnumbered address and preferred source address for the inet family.

```
[edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit" family inet]
user@host# set unnumbered-address lo0.0 preferred-source-address 100.100.100.1
```

Configuring a Dynamic Profile to Dynamically Create IP Demux Interfaces

CLI Quick Configuration

To quickly configure the dynamic profile used to dynamically create IP demux interfaces in the example:

```
[edit]
set dynamic-profiles DHCP-IP-Demux interfaces demux0 unit $junos-interface-unit
  proxy-arp
set dynamic-profiles DHCP-IP-Demux interfaces demux0 unit $junos-interface-unit
  demux-options underlying-interface $junos-underlying-interface
set dynamic-profiles DHCP-IP-Demux interfaces demux0 unit $junos-interface-unit
  family inet demux-source $junos-subscriber-ip-address
set dynamic-profiles DHCP-IP-Demux interfaces demux0 unit $junos-interface-unit
  family inet unnumbered-address lo0.0 preferred-source-address 100.100.100.1
```

Step-by-Step Procedure

To configure a dynamic profile that automatically creates IP demux interfaces:

1. Create a dynamic profile for dynamically creating IP demux interfaces.

```
[edit]
user@host# edit dynamic-profiles DHCP-IP-Demux
```

2. Specify that the dynamic profile use the demux0 interface.

```
[edit dynamic-profiles DHCP-IP-Demux]
user@host# edit interfaces demux0
```

3. Specify that the dynamic profile apply the interface unit value to the dynamic IP demux interfaces.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0]
user@host# edit unit $junos-interface-unit
```

4. (Optional) Configure the router to respond to any ARP request, as long as the router has an active route to the target address of the ARP request.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set proxy-arp
```

5. Specify the logical underlying interface for the dynamic IP demux interfaces.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set demux-options underlying-interface $junos-underlying-interface
```

6. Specify the protocol family information for the dynamic IP demux interfaces.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# edit family inet
```

7. Specify the demux source address is obtained from the incoming subscriber IP address.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0 unit "$junos-interface-unit"
  family inet]
user@host# set demux-source $junos-subscriber-ip-address
```

8. Specify the loopback interface as the unnumbered address and the demux interface IP address as the preferred source address for the dynamic IP demux interfaces.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0 unit "$junos-interface-unit"
  family inet]
user@host# set unnumbered-address lo0.0 preferred-source-address 100.100.100.1
```

Verification

- Subscriber Verification on page 800
- Interface Verification on page 800

Subscriber Verification

Purpose View subscriber information on the router.

Action

- To display dynamic subscriber information:

```
user@host# show subscribers detail
```

Interface Verification

Purpose View interface-specific information on the router.

Action

- To display interface-specific output:

```
user@host# show interfaces interface-name
```

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring a Basic Dynamic Profile on page 673
- Configuring Predefined Dynamic Variables in Dynamic Profiles on page 675

 Copyright © 2013, Juniper Networks, Inc.

- Dynamic 802.1Q VLAN Overview on page 701
- Configuring VLAN Dynamic Profiles on page 714
- *Demultiplexing Interface Overview*

Example: Concurrent Configuration of Dynamic DHCP IP Demux and PPPoE Demux Interfaces over the Same VLAN Demux Interface

This example shows how to configure both dynamic DHCP IP demux and PPPoE demux interfaces over the same dynamic VLAN demux interface. The example provides an IPv4 configuration. However, you can also configure concurrent IP over Ethernet/DHCP and PPPoE interfaces over the same VLAN interface using IPv6 addressing.

NOTE: You can also configure dynamic IP over Ethernet/DHCP and PPPoE interfaces concurrently over the same static VLAN interface. For information on how to configure static VLAN interfaces, see the *Ethernet Interfaces*.

- Requirements on page 801
- Overview on page 801
- Configuration on page 801
- Verification on page 810

Requirements

Before you begin, make sure to configure either DHCP Relay or DHCP Local Server. You can find information about configuring either of these components in the *Junos OS Subscriber Management and Services Library*.

Overview

With the introduction of the **family pppoe** statement, PPPoE is no longer treated as an exclusive encapsulation configuration and you can configure VLAN interfaces with multiple protocol interface stacks. For example, you can configure IP over Ethernet/DHCP and PPPoE interfaces concurrently over a single VLAN interface.

Configuration

- Preparing a Subscriber Access Interface on page 802
- Preparing the Loopback Interface on page 804
- Configuring a Dynamic Profile to Create Dynamic Single-Tagged VLANs on page 804
- Configuring a Dynamic Profile to Create Dynamic Dual-Tagged VLANs on page 806
- Configuring a Dynamic Profile to Create Dynamic IP Demux Interfaces on page 808
- Configuring a Dynamic Profile to Create Dynamic PPPoE Interfaces on page 809

Copyright © 2013, Juniper Networks, Inc.

Preparing a Subscriber Access Interface

CLI Quick Configuration

To quickly configure the aggregated Ethernet interface over which subscribers access the router:

```
[edit]
set chassis aggregated-devices ethernet device-count 1
set interfaces ge-5/0/9 gigether-options 802.3ad ae0
set interfaces ge-5/1/9 gigether-options 802.3ad ae0
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 auto-configure vlan-ranges dynamic-profile Auto-VLAN-Demux accept
   any
set interfaces ae0 auto-configure vlan-ranges dynamic-profile Auto-VLAN-Demux ranges
   ranges 1000-1500
set interfaces ae0 auto-configure stacked-vlan-ranges dynamic-profile
   Auto-Stacked-VLAN-Demux accept any
set interfaces ae0 auto-configure stacked-vlan-ranges dynamic-profile
   Auto-Stacked-VLAN-Demux ranges 1501-2000,any
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp link-protection
```

Step-by-Step Procedure

When configuring multiple protocol interface stacks concurrently over the same VLAN interface, you must configure physical interfaces over which DHCP or PPPoE clients initially access the router. We recommend that you specify the same VLAN tagging for the interface that you expect from incoming clients. This example uses flexible VLAN tagging to simultaneously support transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port.

To automatically create dynamic VLANs, the interface must also include the VLAN range type (single or stacked), dynamic profile reference, and any specific ranges you want the VLANs to use.

To configure a physical interface for subscriber access:

1. Access the physical interface over which you want subscribers to initially access the router.

   ```
   [edit]
   user@host# edit interfaces ge-5/0/9
   ```

2. Specify the aggregated Ethernet interface to which the physical interface belongs.

   ```
   [edit interfaces ge-5/0/9]
   user@host# set gigether-options 802.3ad ae0
   ```

3. Repeat Step 1 and Step 2for each interface you want to assign to the aggregated Ethernet bundle.

   ```
   [edit]
   user@host# set interfaces ge-5/1/9 gigether-options 802.3ad ae0
   ```

4. Access the aggregated Ethernet interface.

   ```
   [edit]
   user@host# edit interfaces ae0
   ```

5. Specify the VLAN tagging that you want the aggregated Ethernet interfaces to use.

Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces ae0]
user@host# set flexible-vlan-tagging
```

6. Edit the **auto-configure** stanza to automatically configure VLANs.

```
[edit interfaces ae0]
user@host# edit auto-configure
```

7. Edit the **vlan-ranges** stanza for single-tagged VLANs.

```
[edit interfaces ae0 auto-configure]
user@host# edit vlan-ranges
```

8. Specify the dynamic VLAN profile that you want the interface to use for dynamically creating single-tagged VLANs.

```
[edit interfaces ae0 auto-configure vlan-ranges]
user@host# edit dynamic-profile Auto-VLAN-Demux
```

9. Specify what VLAN Ethernet packet type the VLAN profile accepts.

```
[edit interfaces ae0 auto-configure vlan-ranges dynamic-profile Auto-VLAN-Demux]
user@host# set accept any
```

10. Specify the VLAN ranges that you want the dynamic profile to use. The following example specifies a lower VLAN ID limit of 1000 and an upper VLAN ID limit of 1500.

```
[edit interfaces ae0 auto-configure vlan-ranges dynamic-profile Auto-VLAN-Demux]
user@host# set ranges 1000-1500
```

11. Edit the **stacked-vlan-ranges** stanza for the dual-tagged VLANs.

```
[edit interfaces ae0 auto-configure]
user@host# edit stacked-vlan-ranges
```

12. Specify the dynamic VLAN profile that you want the interface to use for dynamically creating dual-tagged VLANs.

```
[edit interfaces ae0 auto-configure stacked-vlan-ranges]
user@host# edit dynamic-profile Auto-Stacked-VLAN-Demux
```

13. Specify what VLAN Ethernet packet type the stacked VLAN profile accepts.

```
[edit interfaces ae0 auto-configure stacked-vlan-ranges dynamic-profile
  Auto-Stacked-VLAN-Demux]
user@host# set accept any
```

14. Specify the outer and inner stacked VLAN ranges that you want the dynamic profile to use. The following example specifies an outer stacked VLAN ID range from 1501 through 2000 (to avoid overlapping VLAN IDs with single-tag VLANs) and an inner stacked VLAN ID range of any (enabling a range from 1 through 4094 for the inner stacked VLAN ID).

```
[edit interfaces ge-5/0/9 auto-configure stacked-vlan-ranges dynamic-profile
  Auto-Stacked-VLAN-Demux]
user@host# set ranges 1501-2000,any
```

15. (Optional) Activate the transmission of LACP packets on the aggregated Ethernet interfaces.

```
[edit interfaces ae0]
user@host# set aggregated-ether-options lacp active
```

16. Specify that the aggregated Ethernet interfaces use link protection.

```
[edit interfaces ae0]
user@host# set aggregated-ether-options link-protection
```

Preparing the Loopback Interface

CLI Quick Configuration

To quickly configure the required loopback interface for this example:

```
[edit]
set interfaces lo0.0 unit 0 family inet address 100.100.100.1/32
```

Step-by-Step Procedure

You must configure a loopback interface for use as the unnumbered address and preferred source address for dynamically created interfaces.

To configure the required loopback interface for this example:

1. Configure a loopback interface.

```
[edit]
user@host# edit interfaces lo0.0
```

2. Specify that the loopback interface accept inet packets.

```
[edit interfaces lo0 unit 0]
user@host# edit family inet
```

3. Specify the IP address for the loopback interface.

```
[edit interfaces lo0 unit 0 family inet]
user@host# set address 100.100.100.1/32
```

Configuring a Dynamic Profile to Create Dynamic Single-Tagged VLANs

CLI Quick Configuration

To quickly configure the dynamic profile used to dynamically create single-tagged VLANs in the example:

```
[edit]
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
   demux-source inet
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
   proxy-arp
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
   vlan-id $junos-vlan-id
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
   demux options underlying-interface $junos-interface-ifd-name
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
   family inet unnumbered-address lo0.0 preferred source-address 100.100.100.1
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
   family pppoe duplicate-protection
set dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit $junos-interface-unit
   family pppoe dynamic-profile PPP-Base-PAP
```


Copyright © 2013, Juniper Networks, Inc.

Step-by-Step Procedure

For both dynamic DHCP IP demux and dynamic PPPoE interfaces to reside concurrently on a single-tagged VLAN interface, the VLAN interface must first exist.

To configure a dynamic profile that automatically creates VLAN interfaces:

1. Create a dynamic profile for automatically creating VLAN interfaces.

   ```
   [edit]
   user@host# edit dynamic-profiles Auto-VLAN-Demux
   ```

2. Specify that the dynamic VLAN profile use the demux interface.

   ```
   [edit dynamic-profiles “Auto-VLAN-Demux”]
   user@host# edit interfaces demux0
   ```

3. Specify that the dynamic profile apply the demux interface unit value to the dynamic VLANs.

   ```
   [edit dynamic-profiles Auto-VLAN-Demux interfaces demux0]
   user@host# edit unit $junos-interface-unit
   ```

4. Specify that the demux source accept IPv4 (inet) packets.

   ```
   [edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
     “$junos-interface-unit”]
   user@host# set demux-source inet
   ```

5. (Optional) Specify that each dynamically created interface respond to any ARP request, as long as an active route exists to the target address of the ARP request.

   ```
   [edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
     “$junos-interface-unit”]
   user@host# set proxy-arp
   ```

6. Specify that VLAN IDs are dynamically created.

   ```
   [edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
     “$junos-interface-unit”]
   user@host# set vlan-id $junos-vlan-id
   ```

7. Specify the logical underlying interface for the dynamic VLANs.

   ```
   [edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
     “$junos-interface-unit”]
   user@host# set demux-options underlying-interface $junos-interface-ifd-name
   ```

8. Specify that the VLAN demux interface can accept inet family packets for IP over Ethernet/DHCP subscribers.

   ```
   [edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
     “$junos-interface-unit”]
   user@host# edit family inet
   ```

9. Specify the loopback address as the unnumbered address and preferred source address for the inet family.

   ```
   [edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
     “$junos-interface-unit” family inet]
   user@host# set unnumbered-address lo0.0 preferred-source-address 100.100.100.1
   ```

10. Specify that the VLAN demux interface can accept pppoe family packets for PPPoE subscribers.

```
[edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# edit family pppoe
```

11. Prevent multiple PPPoE sessions from being created for the same PPPoE subscriber on the same VLAN interface.

```
[edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit" family pppoe]
user@host# set duplicate-protection
```

12. Apply the dynamic PPP interface profile to any dynamic PPP interfaces.

```
[edit dynamic-profiles Auto-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit" family pppoe]
user@host# set dynamic-profile PPP-Base-PAP
```

Configuring a Dynamic Profile to Create Dynamic Dual-Tagged VLANs

CLI Quick Configuration

To quickly configure the dynamic profile used to dynamically create stacked/dual-tagged VLANs in the example:

```
[edit]
set dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  $junos-interface-unit demux-source inet
set dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  $junos-interface-unit proxy-arp
set dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  $junos-interface-unit vlan-tags outer $junos-stacked-vlan-id
set dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  $junos-interface-unit vlan-tags inner $junos-vlan-id
set dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  $junos-interface-unit demux options underlying-interface $junos-interface-ifd-name
set dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  $junos-interface-unit family inet unnumbered-address lo0.0 preferred source-address
  100.100.100.1
set dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  $junos-interface-unit family pppoe duplicate-protection
set dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  $junos-interface-unit family pppoe dynamic-profile PPP-Base-PAP
```

Step-by-Step Procedure

For both dynamic DHCP IP demux and dynamic PPPoE interfaces to reside concurrently on a VLAN interface, the VLAN interface must first exist.

To configure a dynamic profile that automatically creates stacked/dual-tagged VLAN interfaces:

1. Create a dynamic profile for automatically creating VLAN interfaces.

```
[edit]
user@host# edit dynamic-profiles Auto-Stacked-VLAN-Demux
```

2. Specify that the dynamic VLAN profile use the demux interface.

```
[edit dynamic-profiles "Auto-Stacked-VLAN-Demux"]
user@host# edit interfaces demux0
```


Copyright © 2013, Juniper Networks, Inc.

3. Specify that the dynamic profile apply the demux interface unit value to the dynamic VLANs.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0]
user@host# edit unit $junos-interface-unit
```

4. Specify that the demux source accept IPv4 (inet) packets.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set demux-source inet
```

5. (Optional) Specify that each dynamically created interface respond to any ARP request, as long as an active route exists to the target address of the ARP request.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set proxy-arp
```

6. Specify that the outer VLAN ID is dynamically created.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set vlan-id -tags outer $junos-stacked-vlan-id
```

7. Specify that the inner VLAN ID is dynamically created.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set vlan-id -tags inner $junos-vlan-id
```

8. Specify the logical underlying interface for the dynamic VLANs.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set demux-options underlying-interface $junos-interface-ifd-name
```

9. Specify that the VLAN demux interface can accept inet family packets for IP over Ethernet/DHCP subscribers.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# edit family inet
```

10. Specify the loopback address as the unnumbered address and preferred source address for the inet family.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit" family inet]
user@host# set unnumbered-address lo0.0 preferred-source-address 100.100.100.1
```

11. Specify that the VLAN demux interface can accept pppoe family packets for PPPoE subscribers.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# edit family pppoe
```

12. Prevent the activation of another dynamic PPPoE logical interface on the same demux underlying interface.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit" family pppoe]
user@host# set duplicate-protection
```

13. Apply the dynamic PPP interface profile to any dynamic PPP interfaces.

```
[edit dynamic-profiles Auto-Stacked-VLAN-Demux interfaces demux0 unit
  "$junos-interface-unit" family pppoe]
user@host# set dynamic-profile PPP-Base-PAP
```

Configuring a Dynamic Profile to Create Dynamic IP Demux Interfaces

CLI Quick Configuration

To quickly configure the dynamic profile used to dynamically create DHCP IP demux interfaces in the example:

```
[edit]
set dynamic-profiles DHCP-IP-Demux interfaces demux0 unit $junos-interface-unit
  proxy-arp
set dynamic-profiles DHCP-IP-Demux interfaces demux0 unit $junos-interface-unit
  demux-options underlying-interface $junos-underlying-interface
set dynamic-profiles DHCP-IP-Demux interfaces demux0 unit $junos-interface-unit
  family inet demux-source $junos-subscriber-ip-address
set dynamic-profiles DHCP-IP-Demux interfaces demux0 unit $junos-interface-unit
  family inet unnumbered-address lo0.0 preferred-source-address 100.100.100.1
```

Step-by-Step Procedure

To configure a dynamic profile that automatically creates IP demux interfaces:

1. Create a dynamic profile for dynamically creating IP demux interfaces.

```
[edit]
user@host# edit dynamic-profiles DHCP-IP-Demux
```

2. Specify that the dynamic profile use the demux0 interface.

```
[edit dynamic-profiles DHCP-IP-Demux]
user@host# edit interfaces demux0
```

3. Specify that the dynamic profile apply the interface unit value to the dynamic PPPoE interfaces.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0]
user@host# edit unit $junos-interface-unit
```

4. (Optional) Configure the router to respond to any ARP request, as long as the router has an active route to the target address of the ARP request.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set proxy-arp
```

5. Specify the logical underlying interface for the dynamic IP demux interfaces.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0 unit
  "$junos-interface-unit"]
user@host# set demux-options underlying-interface $junos-underlying-interface
```

6. Specify the protocol family information for the dynamic IP demux interfaces.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0 unit
  "$junos-interface-unit"]
```


Copyright © 2013, Juniper Networks, Inc.

```
user@host# edit family inet
```

7. Specify the demux source address is obtained from the incoming subscriber IP address.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0 unit "$junos-interface-unit"
  family inet]
user@host# set demux-source $junos-subscriber-ip-address
```

8. Specify the loopback interface as the unnumbered address and the demux interface IP address as the preferred source address for the dynamic IP demux interfaces.

```
[edit dynamic-profiles DHCP-IP-Demux interfaces demux0 unit "$junos-interface-unit"
  family inet]
user@host# set unnumbered-address lo0.0 preferred-source-address 100.100.100.1
```

Configuring a Dynamic Profile to Create Dynamic PPPoE Interfaces

CLI Quick Configuration

To quickly configure the dynamic profile used to dynamically create PPPoE interfaces in the example:

```
[edit]
set dynamic-profiles PPP-Base-PAP interfaces pp0 unit $junos-interface-unit ppp-options
  pap
set dynamic-profiles PPP-Base-PAP interfaces pp0 unit $junos-interface-unit
  pppoe-options underlying-interface $junos-underlying-interface server
set dynamic-profiles PPP-Base-PAP interfaces pp0 unit $junos-interface-unit
  no-keepalives
set dynamic-profiles PPP-Base-PAP interfaces pp0 unit $junos-interface-unit family inet
  unnumbered-address lo0.0
```

Step-by-Step Procedure

1. Create a dynamic profile for automatically creating PPPoE interfaces.

```
[edit]
user@host# edit dynamic-profiles PPP-Base-PAP
```

2. Specify that the dynamic PPPoE profile use the pp0 interface.

```
[edit dynamic-profiles PPP-Base-PAP]
user@host# edit interfaces pp0
```

3. Specify that the dynamic profile apply the interface unit value to the dynamic PPPoE interfaces.

```
[edit dynamic-profiles PPP-Base-PAP interfaces pp0]
user@host# edit unit $junos-interface-unit
```

4. Specify that dynamically created PPPoE interfaces use PAP authentication.

```
[edit dynamic-profiles PPP-Base-PAP interfaces pp0 unit "$junos-interface-unit"]
user@host# set ppp-options pap
```

5. Specify the logical underlying interface for the dynamic PPPoE interfaces.

```
user@host# set ppp-options pap
```

```
[edit dynamic-profiles PPP-Base-PAP interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options underlying-interface $junos-underlying-interface
```

6. Specify that the router act as a PPPoE server.

```
[edit dynamic-profiles PPP-Base-PAP interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options server
```

7. (Optional) Disable the sending of keepalive messages on the dynamic PPPoE interfaces.

```
[edit dynamic-profiles PPP-Base-PAP interfaces pp0 unit "$junos-interface-unit"]
user@host# set no-keepalives
```

8. Specify the protocol family information for the dynamic PPPoE interfaces.

```
[edit dynamic-profiles PPP-Base-PAP interfaces pp0 unit "$junos-interface-unit"]
user@host# edit family inet
```

9. Specify the loopback interface as the unnumbered address for the dynamic PPPoE interfaces.

```
[edit dynamic-profiles PPP-Base-PAP interfaces pp0 unit "$junos-interface-unit"]
user@host# set unnumbered-address lo0.0
```

Verification

- Subscriber Verification on page 810
- Interface Verification on page 810

Subscriber Verification

Purpose View subscriber information on the router.

Action

- To display dynamic subscriber information:

```
user@host# show subscribers detail
```

Interface Verification

Purpose View interface-specific information on the router.

Action

- To display interface-specific output:

```
user@host# show interfaces interface-name
```

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring a Basic Dynamic Profile on page 673
- Configuring Predefined Dynamic Variables in Dynamic Profiles on page 675
- Dynamic 802.1Q VLAN Overview on page 701
- Configuring VLAN Dynamic Profiles on page 714
- *Demultiplexing Interface Overview*
- Configuring the PPPoE Family for an Underlying Interface on page 831

 Copyright © 2013, Juniper Networks, Inc.

Example: Configuring CoS on Static LSQ MLPPP Bundle Interfaces

This example shows how to configure dynamic subscriber services on MLPPP bundle interfaces. The MLPPP bundles must be configured on link services intelligent queuing (IQ) (LSQ) interfaces. The MLPPP interfaces must be statically configured.

To configure dynamic subscriber services on static LSQ MLPPP bundle interfaces:

1. Configure class of service features for the LSQ interfaces.

```
[edit]
class-of-service
  classifiers {
    inet-precedence inet_classifier {
      forwarding-class best-effort {
        loss-priority low code-points 000;
      }
      forwarding-class expedited-forwarding {
        loss-priority low code-points 011;
      }
      forwarding-class assured-forwarding {
        loss-priority low code-points 100;
      }
    }
  }
  fragmentation-maps {
    sample-fragmap {
      forwarding-class {
        best-effort {
          fragment-threshold 1000;
          multilink-class 1:
        }
        assured-forwarding {
          fragment-threshold 1000;
          multilink-class 2:
        }
        expedited-forwarding {
          multilink-class 3:
        }
      }
    }
  }
  forwarding-classes {
    queue 0 best-effort;
    queue 1 expedited-forwarding;
    queue 2 assured-forwarding;
  }
  # traffic classifiers are statically defined
    network traffic interface{
      classifiers {
        inet-precedence inet_classifier;
      }
    }
  scheduler-maps {
```

```
        allthree {
            forwarding-class best-effort scheduler be-scheduler;
            forwarding-class expedited-forwarding scheduler hiprior-sched;
            forwarding-class assured-forwarding scheduler vpn-sched;
        }
    }
    schedulers {
        be-scheduler {
            transmit-rate percent 30;
            priority low;
        }
        hiprior-scheduler {
            transmit-rate percent 40;
            priority strict-high;
        }
        vpn-sched {
            transmit-rate percent 30;
            medium-high;
        }
    }
}
```

2. Configure the MLPPP bundle interfaces and the LSQ interfaces.

```
[edit interfaces]
t1-3/1/0:1:1 {
    keepalives interval 600;
    encapsulation ppp;
    unit 0 {
        ppp-options {
            lcp-restart-timer 5000;
        }
        family mlppp {
            bundle lsq-3/3/0.0;
        }
    }
}
t1-3/1/0:1:2 {
    keepalives interval 600;
    encapsulation ppp;
    unit 0 {
        ppp-options {
            lcp-restart-timer 5000;
        }
        family mlppp {
            bundle lsq-3/3/0.0;
        }
    }
}
lsq-3/3/0 {
    unit 0 {
        encapsulation multilink-ppp;
        multilink-max-classes 4;
        ppp-options {
            ncp-restart-timer 10000;
            dynamic-profile mlppp-profile;
```

 Copyright © 2013, Juniper Networks, Inc.

```
        }
        family inet {
          address 192.168.1.1/32 {
            destination 192.168.25.45;
          }
        }
      }
    }
```

3. Configure the dynamic profile that is applied to the MLPPP bundle interfaces.

```
[edit]
dynamic-profiles {
  mlppp-profile {
    interfaces {
      "$junos-interface-ifd-name" {
        unit junos-underlying-interface-unit" {
          family inet {
            filter {
              input "$junos-input-filter";
              output "$junos-output-filter";
            }
          }
        }
      }
    }
    class-of-service {
      interfaces {
        "$junos-interface-ifd-name" {
          unit junos-underlying-interface-unit" {
            output-traffic-control-profile tcp1;
            fragmentation-map sample-fragmap;
          }
        }
      }
      traffic-control-profiles {
        tcp1 {
          scheduler-map "junos-cos-scheduler-map";
          shaping-rate "$junos-cos-shaping-rate";
          guaranteed-rate "$junos-cos-guaranteed-rate";
          delay-buffer-rate "$junos-cos-delay-buffer-rate";
        }
      }
      scheduler-maps {
        data_smap {
          forwarding-class be scheduler data_sch;
        }
      }
      schedulers {
        be_sch {
          ...
        }
      }
    }
  }
}
```

Related Documentation

- For hardware requirements, see Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces on page 374
- For information about MLPPP and LSQ interfaces, see *Link Services IQ Interfaces Configuration* in the *Junos OS Services Interfaces Library for Routing Devices*.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 43

Subscriber Interfaces over Aggregated Ethernet Overview

- Static and Dynamic VLAN Subscriber Interfaces over Aggregated Ethernet Overview on page 815
- Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview on page 816
- Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 818

Static and Dynamic VLAN Subscriber Interfaces over Aggregated Ethernet Overview

You can configure a subscriber interface represented by a static virtual LAN (VLAN) stacked on a two-link aggregated Ethernet logical interface. You must configure the aggregated Ethernet logical interface on Enhanced Queuing Dense Port Concentrators (EQ DPCs) or MPC/MIC interfaces in MX Series 3D Universal Edge Routers.

A static or dynamic VLAN subscriber interface over aggregated Ethernet can also support one-to-one active/backup link redundancy, depending on how you configure the underlying aggregated Ethernet interface.

To configure a static or dynamic VLAN subscriber interface over aggregated Ethernet, make sure you understand the following concepts.

- Guidelines for Configuring an Aggregated Ethernet Logical Interface to Support a Static or Dynamic VLAN Subscriber Interface on page 815

Guidelines for Configuring an Aggregated Ethernet Logical Interface to Support a Static or Dynamic VLAN Subscriber Interface

The following guidelines for configuring an aggregated Ethernet logical interface also apply to configuring a static or dynamic VLAN subscriber interface stacked on a two-link aggregated Ethernet logical interface:

- If you need to support one-to-one active/backup link redundancy, configure the aggregated Ethernet interface in link protection mode, which requires that the two underlying physical interfaces be designated as primary and backup links.

Copyright © 2013, Juniper Networks, Inc.

- In addition, if you need to support one-to-one active/backup link redundancy at the DPC or MPC level, configure the aggregated Ethernet interface on physical interfaces that reside on different EQ DPCs or MPCs.

NOTE: **One-to-one active/backup DPC redundancy is also supported with firewall filters and policy filters for static non-VLAN interfaces configured on an aggregated Ethernet logical interfaces, provided LACP is not active.**

Related Documentation

- Static Subscriber Interfaces and VLAN Overview on page 758
- Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet on page 824
- Example: Configuring a Static Subscriber Interface on a VLAN Interface over Aggregated Ethernet on page 835
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for Subscriber Access Overview on page 949

Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview

You can configure a subscriber interface using a static or dynamic demux interface stacked on an aggregated Ethernet logical interface. Subscriber interfaces on static or dynamic demux interfaces can be used to identify specific subscribers (authenticated users) in an access network or to separate individual circuits. A subscriber interface on a static or dynamic demux interface over aggregated Ethernet can support one-to-one active/backup link redundancy or traffic load balancing, depending on how you configure the underlying aggregated Ethernet interface.

To configure a static or dynamic demux subscriber interface over aggregated Ethernet, make sure you understand the following concepts:

- Options for Aggregated Ethernet Logical Interfaces That Support Demux Subscriber Interfaces on page 816
- Hardware Requirements with Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet on page 817
- Features Supported with Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet on page 817

Options for Aggregated Ethernet Logical Interfaces That Support Demux Subscriber Interfaces

Traffic forwarding through a demux logical interface is dependent on the configuration of the underlying interface. Using an aggregated Ethernet interface as the underlying interface for a static or dynamic demux subscriber interface provides you with the following options:

- **1:1 Active/Backup Link Redundancy**—If you need to support one-to-one active/backup link redundancy, configure the aggregated Ethernet interface in link protection mode,

 Copyright © 2013, Juniper Networks, Inc.

which requires that two underlying physical interfaces be designated as primary and backup links. In addition, if you need to support one-to-one active/backup link redundancy at the line card level, configure the aggregated Ethernet interface on physical interfaces that reside either on different EQ DPCs or on different MPCs. When using LACP link protection, you can configure only two member links to an aggregated Ethernet interface: one active and one standby.

- **Load Balancing**—You can configure load balancing instead of 1:1 active/backup link redundancy. The Junos OS implementation of the IEEE 802.3ad standard balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet.

  By default, the system supports hash-based distribution in load balancing scenarios. In this model, traffic for a logical interface can be distributed over multiple links in the aggregated Ethernet interface. If distribution flows are not even, egress CoS scheduling can be inaccurate. In addition, scheduler resources are required on every link of the aggregated Ethernet interface.

  Targeted distribution enables you to target the egress traffic for IP and VLAN demux subscribers on a single member link, using a single scheduler resource. The system distributes the subscriber interfaces equally among the member links.

Hardware Requirements with Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet

IP demux subscriber interfaces over aggregated Ethernet interfaces are supported on EQ DPCs on MX Series routers.

VLAN demux subscriber interfaces over aggregated Ethernet interfaces are supported on MX Series routers that only have MPCs installed. If the router has other line cards in addition to MPCs, the CLI accepts the configuration but errors are reported when the subscriber interfaces are brought up.

Features Supported with Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet

Table 69 on page 818 lists key subscriber access features supported with static or dynamic demux subscriber interfaces, organized by type of underlying interface:

- Aggregated Ethernet
- Non-aggregated Ethernet (Gigabit Ethernet, Fast Ethernet, or 10-Gigabit Ethernet)

There are no feature limitations specific to demultiplexing. Instead, demux interfaces over aggregated Ethernet are subject to the same scaling and configuration limitations inherent to aggregated Ethernet logical interfaces.

Copyright © 2013, Juniper Networks, Inc.

Table 69: Features Supported with Static or Dynamic Demux Subscriber Interfaces

Feature	Static or Dynamic Demux Subscriber Interface	
	Aggregated Ethernet Underlying Interface	Non-aggregated Underlying Logical Interface
Protocol family support	IPv4, IPv6, and PPPoE	IPv4, IPv6, and PPPoE
Per-subscriber firewall filtering and statistics	Supported	Supported
Hierarchical CoS	Supported	Supported
Per-subscriber CoS parameters within the **[edit dynamic-profiles *profile-name* class-of-service]** hierarchy	Supported	Supported
Per-subscriber IGMP configuration within the **[edit dynamic-profiles *profile-name* protocols]** hierarchy NOTE: IP demux interfaces must use OIF mapping. See *Example: Configuring Multicast with Subscriber VLANs* in the *Multicast Protocols Feature Guide for Routing Devices* for additional information.	Yes	Yes

Related Documentation

- Subscriber Interfaces and Demultiplexing Overview on page 759
- Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 818
- Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet on page 825
- Configuring the PPPoE Family for an Underlying Interface on page 831
- Example: Configuring a Static Subscriber Interface on an IP Demux Interface over Aggregated Ethernet on page 838
- For more information about aggregated Ethernet interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

Distribution of Demux Subscribers in an Aggregated Ethernet Interface

This topic describes the distribution options available for demux subscriber interfaces over aggregated Ethernet.

Distribution Models

By default, the system supports hash-based distribution for all subscriber interface types in an aggregated Ethernet bundle configured without link protection. In this model, traffic for a logical interface can be distributed over multiple links in the bundle. This model is desirable when there are many flows through the logical interface and you need to load balance those flows.

Copyright © 2013, Juniper Networks, Inc.

Note that if the distribution flows are not even, egress CoS scheduling can be inaccurate. In addition, scheduler resources are required on every link of the aggregated Ethernet interface. For example, if subscriber traffic is allocated 10 MB for a triple-play service over four links in a bundle, each of the links could receive 2.5 MB of traffic. High-density services such as video could be limited by the bandwidth on one of the links.

Targeted distribution enables you to target the egress traffic for an IP or VLAN demux subscriber on a single member link, using a single scheduler resource. To achieve load balancing over the member links, the system distributes the subscriber interfaces equally among the links. This enables the subscriber that is allocated 10 MB to be accurately scheduled as the traffic flows through.

Sample Targeted Distribution Topology

Figure 15 on page 819 displays a sample targeted distribution of subscriber traffic across links in an aggregated Ethernet interface. A primary and backup link is allocated for each subscriber.

Figure 15: Targeted Subscriber Links

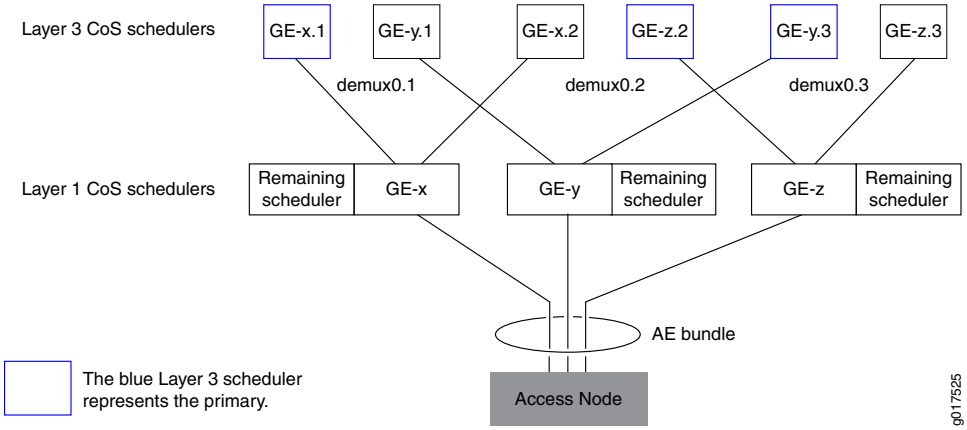

For example, if link **GE-x** went down, subscriber 1 can begin forwarding over the backup, which is link **Ge-y**. When link **GE-y** comes back up, subscriber 1 switches back to its primary link, **GE-x**.

In the event that both **GE-x** and **GE-y** go down, subscriber 3 starts forwarding through its backup, **GE-z**. Subscriber 1 will have lost its primary and backup links, and will also begin forwarding out the **GE-z** link. A new level 3 scheduler is assigned for this subscriber on link **GE-z**. If there is a momentary lapse between the time that a new scheduler is allocated and forwarding switches to **GE-z**, the traffic will be forwarding through to the remaining scheduler. Subscriber 2 continues to forward through its primarily link, **GE-z**.

Redundancy and Redistribution Mechanisms

Two types of redundancy are available in the targeted distribution model: link redundancy and module redundancy.

By default, an aggregated Ethernet interface is enabled with link redundancy. Backup links for a subscriber are chosen based on the link with the least number of subscribers, which provides redundancy if a link fails.

The module redundancy option enables you to provide redundancy if a module or a link fails. Backup links for a subscriber are chosen on a different DPC or MPC from the primary link, based on the link with the least number of subscribers among the links on different modules. You can enable this for the aggregated Ethernet interface.

When links are removed, affected subscribers are redistributed among the active remaining backup links. When links are added to the system, no automatic redistribution occurs. New subscribers are assigned to the links with the fewest subscribers (which are typically the new links).

Considerations and Best Practices

Keep the following guidelines in mind when configuring targeted distribution for demux subscribers:

- You can manage subscribers with both hash-based and targeted distribution models in the same network. For example, you can allocate subscribers with interface types such as PPPoE with hash-based distribution, and enable demux subscribers with targeted distribution.
- We recommend that you configure module redundancy to protect against module failures. When module redundancy is enabled, you can ensure an even distribution of subscribers if you allocate no more than 50 percent of the links on a single DPC or MPC.
- During normal network operations, the system maintains an even balance of subscribers among the links in a bundle, even as subscribers log in and out. However, if the distribution of a bundle becomes uneven (for example, when a link goes down and new subscribers are logging in), you can perform a manual rebalance of the bundle. In addition, you can configure periodic rebalancing of the bundle with a specific time interval.
- When you anticipate that a link will be down for an extended time, and you want to ensure that backup links are provisioned for all subscribers, we recommend that you remove the failed link from the bundle. This forces the affected subscribers to redistribute to other links.
- We recommend that you apply a remaining traffic-control profile to the logical interface to ensure that minimal scheduling parameters are applied to the remaining subscriber traffic. This provides scheduling for subscribers that do not have schedulers allocated because they have not been configured or they have been over-provisioned, or because of scheduler transitions on multiple link failures.
- If you perform a cold restart on the router when it is forwarding active subscribers, the subscriber interfaces with targeted distribution are assigned to the first links that become available when the system is initializing so forwarding can begin. To rebalance the system following a cold restart, perform a manual rebalance of the bundle. In addition, we recommend that you configure Graceful Routing Engine switchover (GRES) on the router to enable nonstop forwarding during switchover, and avoid performing cold restarts.

 Copyright © 2013, Juniper Networks, Inc.

- To ensure appropriate and predictable targeted distribution, you must configure chassis network services to use **enhanced-ip** mode.
- Unless specifically separated, multicast traffic egresses in parallel with unicast traffic, sharing the CoS hierarchy and aggregated Ethernet flow distribution.

Related Documentation

- Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 828
- Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface on page 829
- Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface on page 829
- For additional information about demux over aggregated Ethernet interfaces, see Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview on page 816

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 44

Configuring Subscriber Interfaces over Aggregated Ethernet

- Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet on page 824
- Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet on page 825
- Configuring a Static or Dynamic VLAN Demux Subscriber Interface over Aggregated Ethernet on page 826
- Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 828
- Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface on page 829
- Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface on page 829
- Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 830
- Configuring the PPPoE Family for an Underlying Interface on page 831
- Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces on page 833
- Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface on page 833

Copyright © 2013, Juniper Networks, Inc.

Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet

You can configure a subscriber link represented by a static virtual LAN (VLAN) stacked on an aggregated Ethernet logical interface.

You can configure subscriber management services such as firewall filters and CoS for this subscriber interface.

To configure a subscriber interface using a static VLAN interface over an aggregated Ethernet logical interface:

1. Configure the aggregated Ethernet interface.
   a. Configure the number of aggregated Ethernet interfaces on the router.

      See *Configuring the Number of Aggregated Ethernet Interfaces on the Device*.

   b. Configure the aggregated Ethernet interface.

      See *Configuring an Aggregated Ethernet Interface*.

   c. (Optional) Configure LACP.

      See *Configuring Aggregated Ethernet LACP*.

   d. (Optional) Configure the minimum number of links.

      See *Configuring Aggregated Ethernet Minimum Links*.

   e. (Optional) Configure the link speed.

      See *Configuring Aggregated Ethernet Link Speed*.

   f. (Optional) Configure the aggregated Ethernet logical interface to support one-to-one active/backup link redundancy or traffic load balancing.

      See *Configuring Aggregated Ethernet Link Protection*.

      NOTE: **Link protection is required if you want to configure hierarchical CoS on the aggregated Ethernet interface. For more information, see** "Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links" on page 1023.

2. Configure the static or dynamic VLAN interface.
   - For static VLAN interfaces, see "Configuring a Subscriber Interface with a Static VLAN Interface" on page 766.
   - For dynamic VLAN interfaces, see "Configuring VLAN Dynamic Profiles" on page 714 and "Configuring VLAN Interfaces to Use Dynamic Profiles" on page 720.
3. Configure subscriber management services on the subscriber interface.

 Copyright © 2013, Juniper Networks, Inc.

- For firewall filters, see "Dynamically Attaching Statically Created Filters for Any Interface Type" on page 1170 or "Dynamically Attaching Statically Created Filters for a Specific Interface Family Type" on page 1169.
- For hierarchical CoS, see "Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links" on page 1023.

Related Documentation

- Static and Dynamic VLAN Subscriber Interfaces over Aggregated Ethernet Overview on page 815
- Example: Configuring a Static Subscriber Interface on a VLAN Interface over Aggregated Ethernet on page 835
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for Subscriber Access Overview on page 949

Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet

You can configure a subscriber interface using a static or dynamic IP demultiplexing (demux) logical interface stacked on an aggregated Ethernet logical interface. Optionally, you can configure the aggregated Ethernet logical interface to support one-to-one active/backup link redundancy or traffic load balancing.

1. Configure the aggregated Ethernet interface.
   a. Configure the number of aggregated Ethernet interfaces on the router.

      See *Configuring the Number of Aggregated Ethernet Interfaces on the Device*.

   b. Configure the aggregated Ethernet interface.

      See *Configuring an Aggregated Ethernet Interface*.

   c. (Optional) Configure LACP.

      See *Configuring Aggregated Ethernet LACP*.

   d. (Optional) Configure the minimum number of links.

      See *Configuring Aggregated Ethernet Minimum Links*.

   e. (Optional) Configure the link speed.

      See *Configuring Aggregated Ethernet Link Speed*.

   f. (Optional) Configure the aggregated Ethernet logical interface to support one-to-one active/backup link redundancy or traffic load balancing.

      For general instructions, see *Configuring Aggregated Ethernet Link Protection*.

NOTE: Link protection is required if you want to configure hierarchical CoS on the aggregated Ethernet interface. For more information, see "Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links" on page 1023.

2. Configure the aggregated Ethernet logical interface as the underlying interface to support the static or dynamic IP demux subscriber interface.

   The aggregated Ethernet interface needs to support demultiplexing of incoming traffic to the Ethernet links based on IPv4 destination or source addresses in the incoming packets. In addition, you must configure the IP address of each link.

   See *Configuring an IP Demux Underlying Interface*.

3. Configure the static or dynamic IP demux interface.

   - For static subscriber interfaces, see "Configuring Static Subscriber Interfaces Using IP Demux Interfaces" on page 766.
   - For dynamic subscriber interfaces, see "Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles" on page 771.

   NOTE: IP demux interfaces currently support only the Internet Protocol version 4 (IPv4) suite (family inet).

4. (Optional) Configure subscriber management services on the subscriber interface.

   - For firewall filters, see "Dynamically Attaching Statically Created Filters for Any Interface Type" on page 1170 or "Dynamically Attaching Statically Created Filters for a Specific Interface Family Type" on page 1169.
   - For hierarchical CoS, see "Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links" on page 1023.

Related Documentation

- Subscriber Interfaces and Demultiplexing Overview on page 759
- Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview on page 816
- Example: Configuring a Static Subscriber Interface on an IP Demux Interface over Aggregated Ethernet on page 838
- Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 828

Configuring a Static or Dynamic VLAN Demux Subscriber Interface over Aggregated Ethernet

You can configure a subscriber interface using a static or dynamic VLAN demultiplexing (demux) logical interface stacked on an aggregated Ethernet physical interface.

 Copyright © 2013, Juniper Networks, Inc.

1. Configure the aggregated Ethernet interface.
   a. Configure the number of aggregated Ethernet interfaces on the router.

      See *Configuring the Number of Aggregated Ethernet Interfaces on the Device*.

   b. Configure the aggregated Ethernet interface.

      See *Configuring an Aggregated Ethernet Interface*.

   c. (Optional) Configure LACP.

      See *Configuring Aggregated Ethernet LACP*.

   d. (Optional) Configure the minimum number of links.

      See *Configuring Aggregated Ethernet Minimum Links*.

   e. (Optional) Configure the link speed.

      See *Configuring Aggregated Ethernet Link Speed*.

   f. (Optional) Configure the aggregated Ethernet logical interface to support one-to-one active/backup link redundancy or traffic load balancing.

      For general instructions, see *Configuring Aggregated Ethernet Link Protection*.

2. Configure the aggregated Ethernet physical interface as the underlying interface to support the static or dynamic VLAN demux subscriber interface.

   The aggregated Ethernet interface needs to support demultiplexing of incoming traffic to the Ethernet links based on the VLAN ID in the incoming packets.

   See *Configuring a VLAN Demux Underlying Interface*.

3. Configure the static or dynamic VLAN demux interface.
   - For static subscriber interfaces, see "Configuring Static Subscriber Interfaces Using VLAN Demux Interfaces" on page 767.
   - For dynamic subscriber interfaces, see "Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles" on page 772.

   NOTE: **VLAN demux interfaces currently support the Internet Protocol version 4 (IPv4) suite (family inet) and the Internet Protocol version 6 (IPv6) suite (family inet6).**

   VLAN demux subscriber interfaces over aggregated Ethernet physical interfaces are supported only for MX Series routers that have only MPCs installed. If the router has other cards in addition to MPCs, the CLI accepts the configuration but errors are reported when the subscriber interfaces are brought up.

4. (Optional) Configure subscriber management services on the subscriber interface.

- For firewall filters, see "Dynamically Attaching Statically Created Filters for Any Interface Type" on page 1170 or "Dynamically Attaching Statically Created Filters for a Specific Interface Family Type" on page 1169.
- For hierarchical CoS, see "Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links" on page 1023.

Related Documentation

- Subscriber Interfaces and Demultiplexing Overview on page 759
- Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview on page 816
- *Associating VLAN IDs to VLAN Demux Interfaces*
- Example: Configuring IPv4 Static VLAN Demux Interfaces over a Gigabit Ethernet Underlying Interface with DHCP Local Server on page 791
- Example: Configuring IPv4 Static VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server on page 843
- Example: Configuring IPv4 Dynamic VLAN Demux Interfaces over a Gigabit Ethernet Underlying Interface with DHCP Local Server on page 793
- Example: Configuring IPv4 Dynamic VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server on page 856

Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces

By default, the system supports hash-based distribution of subscriber traffic in aggregated Ethernet bundles. You can configure the system to target the egress traffic for a subscriber on a single member link, using a single scheduler resource. The system distributes the subscriber interfaces equally among the member links.

To configure targeted distribution:

1. Edit the chassis hierarchy level.

   [edit]
   user@host#**edit chassis**

2. Enable chassis network services for **enhanced-ip** mode.

   [edit chassis]
   user@host#**set network-services enhanced-ip**

3. Access the logical interface.

   [edit]
   user@host#**edit interfaces demux0 unit** ***logical-unit-number***

4. Enable targeted distribution for the interface.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host#**set** targeted-distribution

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 830
- Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 818

Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface

By default, an aggregated Ethernet bundle with targeted distribution is enabled with link redundancy. Backup links for a subscriber are chosen based on the link with the fewest subscribers, which provides redundancy if a link fails.

We recommend that you configure the module redundancy option to provide redundancy if a module or a link fails. Backup links for a subscriber are chosen on a different DPC or MPC from the primary link, based on the link with the fewest subscribers among the links on different modules.

To configure module redundancy for an aggregated Ethernet bundle:

1. Access the aggregated Ethernet bundle for which you want to configure module redundancy.

   edit
   user@host# **edit interfaces ae*x* aggregated-ether-options**

2. Enable module redundancy for the bundle.

   [edit interfaces ae*x* aggregated-ether-options]
   user@host# **logical-interface-fpc-redundancy**

Related Documentation

- Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 828
- Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 818

Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface

In a targeted distribution model, the system allocates demux subscriber interfaces equally among the member links in the aggregated Ethernet interface. When links are removed, affected subscribers are redistributed among the active remaining backup links. When links are added to the system, no automatic redistribution occurs. New subscribers are assigned to the links with the fewest subscribers (which are typically the new links).

During normal network operations, the system maintains an even balance of traffic among the links in a bundle, even as subscribers log in and out. However, if the distribution of a bundle becomes uneven (for example, when a link goes down for a period of time

and new subscribers are logging in), you can perform a manual rebalance of the bundle. In addition, you can configure periodic rebalancing of the bundle with a specific interval.

- Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface on page 830
- Configuring Manual Rebalancing of Subscribers on an Aggregated Ethernet Interface on page 830

Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface

If subscribers are frequently logging in and logging out of your network, you can configure the system to periodically rebalance the links based on a specific time and interval.

To configure periodic rebalancing:

1. Access the aggregated Ethernet interface for which you want to configure periodic rebalancing.

   edit
   user@host# **edit interfaces ae*number* aggregated-ether-options**

2. Configure the rebalancing parameters for the interface, including the time and the interval between rebalancing actions.

   [edit interfaces ae*number* aggregated-ether-options]
   user@host# rebalance-periodic **time** *hour:minute* <**interval** *hours*>

Configuring Manual Rebalancing of Subscribers on an Aggregated Ethernet Interface

To manually rebalance the subscribers among the links in an aggregated Ethernet bundle with targeted distribution:

- Issue the **request interface rebalance** command:

  user@host# **request interface rebalance interface** <*interface-name*>

Related Documentation

- Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 830
- Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 828
- Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 818

Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface

Purpose View the distribution status of subscribers that are targeted to links in an aggregated Ethernet interface.

Action

- To display a summary of the distribution of links on the demux interface:

  user@host> **show interfaces demux0 extensive**

- To display the targeted distribution on a specific aggregated Ethernet interface:

 Copyright © 2013, Juniper Networks, Inc.

```
user@host> show interfaces targeting aex
```

Related Documentation

- Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 828
- Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface on page 829

Configuring the PPPoE Family for an Underlying Interface

You can configure the PPPoE family on an underlying interface as an alternative to configuring PPPoE encapsulation on that interface. You cannot configure both on the same interface. You can configure the same attributes for the PPPoE family as you can for an interface configured with **pppoe-underlying-options**.

Before you begin, configure the underlying interface. When you want to configure PPPoE on an aggregated Ethernet bundle, you must configure the PPPoE family over a VLAN demux interface as an intermediate underlying option. The VLAN demux interface can be static or dynamic.

The following topics describe how to configure basic static and dynamic interfaces:

- Configuring Static Subscriber Interfaces in Dynamic Profiles on page 765
- Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles on page 772

To configure the PPPoE family over an underlying interface:

1. Specify the PPPoE family.

   ```
   [edit interfaces demux0 unit logical-unit-number]
   user@host# set family pppoe
   ```

2. (Optional) Configure an alternative access concentrator name to be used instead of the system name in PPPoE control packets for the dynamic PPPoE subscriber interface.

   ```
   [edit interfaces demux0 unit logical-unit-number family pppoe]
   user@host# set access-concentrator name
   ```

3. (Optional) Configure duplicate protection to prevent the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same MAC address is already active on that interface.

   ```
   [edit interfaces demux0 unit logical-unit-number family pppoe]
   user@host# set duplicate-protection
   ```

4. (Optional) Attach a dynamic profile to determine the properties of the dynamic PPPoE logical interface when it is created.

   ```
   [edit interfaces demux0 unit logical-unit-number family pppoe]
   user@host# set dynamic-profile profile-name
   ```

5. (Optional) Configure the maximum number of concurrent PPPoE sessions that the router can activate on the underlying interface in either of the following ways:

   - To configure the maximum number of concurrent PPPoE sessions on a per-interface basis, from 1 through the platform-specific default for your router, use the **max-sessions** statement:

     [edit interfaces demux0 unit *logical-unit-number* family pppoe]
     user@host# **set** max-sessions *number*

   - To configure the maximum number of concurrent PPPoE sessions on a per-subscriber basis, use the value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks vendor-specific attribute (VSA) [26-143]. By default, the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA takes precedence over the PPPoE maximum session value configured with the **max-sessions** statement.

6. (Optional) Configure the router to ignore the value returned by RADIUS in the Max-Clients-Per-Interface VSA and restore the PPPoE maximum session value on the underlying interface to the value configured in the CLI with the **max-sessions** statement.

   [edit interfaces demux0 unit *logical-unit-number* family pppoe]
   user@host# **set** max-sessions-vsa-ignore

7. (Optional) Enable PPPoE subscriber session lockout on the PPPoE underlying interface in either of the following ways:

   - To configure PPPoE subscriber session lockout with the default lockout period:

     [edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]
     user@host# **set** short-cycle-protection

   - To configure PPPoE subscriber session lockout with a nondefault lockout period:

     [edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]
     user@host# **set** short-cycle-protection **lockout-time-min** *minimum-seconds* **lockout-time-max** *maximum-seconds*

   BEST PRACTICE: **When you configure PPPoE subscriber session lockout, we recommend that you also enable duplicate protection to ensure that the MAC source address for each PPPoE session is unique on the underlying interface.**

8. (Optional) Specify the service name table assigned to the underlying interface.

   [edit interfaces demux0 unit *logical-unit-number* family pppoe]
   user@host# **set** service-name-table *table-name*

Related Documentation

- Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview on page 816
- Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces on page 905

Copyright © 2013, Juniper Networks, Inc.

- Configuring Lockout of PPPoE Subscriber Sessions on page 912
- Example: Configuring a Static PPPoE Subscriber Interface on a Static Underlying VLAN Demux Interface over Aggregated Ethernet on page 865
- Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Underlying VLAN Demux Interface over Aggregated Ethernet on page 870
- Example: Configuring a Dynamic PPPoE Subscriber Interface on a Dynamic Underlying VLAN Demux Interface over Aggregated Ethernet on page 876

Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces

By default, the system supports hash-based distribution of subscriber traffic in aggregated Ethernet bundles. You can configure the system to target the egress traffic for a subscriber on a single member link, using a single scheduler resource. The system distributes the subscriber interfaces equally among the member links.

To configure targeted distribution:

1. Edit the chassis hierarchy level.

   [edit]
   user@host#**edit chassis**

2. Enable chassis network services for **enhanced-ip** mode.

   [edit chassis]
   user@host#**set network-services enhanced-ip**

3. Access the logical interface.

   [edit]
   user@host#**edit interfaces pp0 unit** ***logical-unit-number***

4. Enable targeted distribution for the interface.

   [edit interfaces pp0 unit *logical-unit-number*]
   user@host#**set targeted-distribution**

Related Documentation

- CoS for PPPoE Subscriber Interfaces Overview on page 1014
- Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface on page 833

Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface

Purpose View the distribution status of subscribers that are targeted to links in an aggregated Ethernet interface.

Action

- To display a summary of the distribution of links on the demux interface:

  user@host> **show interfaces pp0 extensive**

- To display the targeted distribution on a specific aggregated Ethernet interface:

```
user@host> show interfaces targeting aex
```

Related Documentation

- CoS for PPPoE Subscriber Interfaces Overview on page 1014
- Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces on page 833

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 45

Subscriber Interfaces over Aggregated Ethernet Examples

- Example: Configuring a Static Subscriber Interface on a VLAN Interface over Aggregated Ethernet on page 835
- Example: Configuring a Static Subscriber Interface on an IP Demux Interface over Aggregated Ethernet on page 838
- Example: Configuring a Static Subscriber Interface on a VLAN Interface over Aggregated Ethernet on page 840
- Example: Configuring IPv4 Static VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server on page 843
- Example: Separating Targeted Multicast Traffic for Demux Subscribers on Aggregated Ethernet Interfaces on page 846
- Example: Configuring IPv4 Dynamic VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server on page 856
- Example: Configuring IPv6 Dynamic VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server on page 859
- Example: Configuring IPv4 Dynamic Stacked VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server on page 862
- Example: Configuring a Static PPPoE Subscriber Interface on a Static Underlying VLAN Demux Interface over Aggregated Ethernet on page 865
- Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Underlying VLAN Demux Interface over Aggregated Ethernet on page 870
- Example: Configuring a Dynamic PPPoE Subscriber Interface on a Dynamic Underlying VLAN Demux Interface over Aggregated Ethernet on page 876

Example: Configuring a Static Subscriber Interface on a VLAN Interface over Aggregated Ethernet

This example shows how you can configure a subscriber interface using a static virtual LAN (VLAN) stacked on a two-link aggregated Ethernet logical interface. In this example, the underlying aggregated Ethernet logical interface is configured for one-to-one active/backup redundancy at the DPC level, and per-subscriber static hierarchical

class-of-service (CoS) is configured by applying CoS parameters at the aggregated Ethernet logical interface.

1. Define the number of aggregated Ethernet interfaces on the router.

   In this example, only one aggregated Ethernet logical interface is configured on the router.

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
    }
  }
}
```

2. Configure **ae0**, a two-link aggregated Ethernet logical interface to serve as the underlying interface for the static VLAN subscriber interface. In order to support hierarchical CoS, the physical ports must be on EQ DPCs in MX Series routers.

   In this example, the LAG bundle is configured for one-to-one active/backup link redundancy. To support link redundancy at the DPC level, the LAG bundle attaches ports from two different EQ DPCs.

```
[edit]
interfaces {
  ge-5/0/3 {
    gigether-options {
      802.3ad {
        ae0;
        primary;
      }
    }
    ge-5/1/2 {
      gigether-options {
        802.3ad {
          ae0;
          backup;
        }
      }
    }
  }
}
```

3. Configure **ae0** to serve as the underlying interface for the static VLAN interface.

```
[edit]
interfaces {
  ae0 {
    hierarchical-scheduler;
    aggregated-ether-options {
      link-protection;
      minimum-links 1;
      link-speed 1g;
      lacp {
        active;
```

 Copyright © 2013, Juniper Networks, Inc.

```
            }
          }
        }
      }
```

4. Configure static traffic-shaping and scheduling parameters.

```
[edit]
class-of-service {
  forwarding-classes { # Associate queue numbers with class names
    queue 0 be;
    queue 1 e;
    queue 2 af;
    queue 3 nc;
  }
  schedulers { # Define output queue properties
    scheduler_be {
      transmit-rate percent 30;
      buffer-size percent 30;
    }
    scheduler_ef {
      transmit-rate percent 40;
      buffer-size percent 40;
    }
    scheduler_af {
      transmit-rate percent 25;
      buffer-size percent 25;
    }
    scheduler_nc {
      transmit-rate percent 5;
      buffer-size percent 5;
    }
  }
  scheduler-maps { # Associate queues with schedulers
    smap_2 {
      forwarding-class be scheduler_be;
      forwarding-class ef scheduler_ef;
      forwarding-class-af scheduler_af;
      forwarding-class-nc scheduler_nc;
    }
  }
}
```

5. Attach static CoS to the physical and logical interfaces of the aggregated Ethernet interface.

   In this example, three traffic control profiles are defined, but only two profiles are applied to the static VLAN subscriber interface over aggregated Ethernet:

   - The **tcp_for_ae_device_pir_500m** profile defines a shaping rate, and it is applied to both of the underlying physical interfaces (**ge-5/0/3** and **ge-5/1/2**).
   - The **tcp-for-ae_smap_video_pir_20m_delay_30m** profile defines a scheduler map, a shaping rate, and a delay buffer rate, and it is applied to one of the logical interfaces on the aggregated Ethernet bundle (**ae0.0**).

```
[edit]
```

```
class-of-service {
  traffic-control-profiles { # Configure traffic shaping and scheduling profiles
    tcp_for_ae_device_pir_500m {
      shaping-rate 20m;
    }
    tcp_for_ae_smap_video_pir_20m_delay_30m {
      scheduler-map smap_video;
      shaping-rate 20m;
      delay-buffer-rate 30m;
    }
    tcp_for_ae_smap_video_cir_50m_delay_75m {
      scheduler-map smap_video;
      guaranteed-rate 50m;
      delay-buffer-rate 75m;
    }
  }
  interfaces { # Apply two traffic-control profiles to the LAG
    ae0 { # Two underlying physical interfaces on separate EQ DPCs
      output-traffic-control-profile tcp-for-ae_device_pir_500m;
      unit 0 { # One of the two logical interfaces on 'ae0'
        output-traffic-control-profile tcp-for-ae_smap_video_pir_20m_delay_30m;
      }
    }
  }
}
```

Related Documentation

- Static and Dynamic VLAN Subscriber Interfaces over Aggregated Ethernet Overview on page 815
- Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet on page 824
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for Subscriber Access Overview on page 949

Example: Configuring a Static Subscriber Interface on an IP Demux Interface over Aggregated Ethernet

This example shows how you can configure a subscriber interface using a static IP demultiplexing (demux) interface stacked on a two-link aggregated Ethernet logical interface. In this example, the underlying aggregated Ethernet logical interface is configured for one-to-one active/backup redundancy at the DPC level.

1. Define the number of aggregated Ethernet interfaces on the router.

   In this example, only one aggregated Ethernet logical interface is configured on the router:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
```

 Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
```

2. Configure **ae0**, a two-link aggregated Ethernet logical interface to serve as the underlying interface for the static IP demux subscriber interface.

   In this example, the LAG bundle is configured for one-to-one active/backup link redundancy. To support link redundancy at the DPC level, the LAG bundle attaches ports from two different EQ DPCs.

```
[edit]
interfaces {
  ge-5/0/3 {
    gigether-options {
      802.3ad {
        ae0;
        primary;
      }
    }
  }
  ge-5/1/2 {
    gigether-options {
      802.3ad {
        ae0;
        backup;
      }
    }
  }
}
```

3. Configure the aggregated Ethernet logical interface with link protection enabled, and specify the logical demultiplexing source family type for both the active and backup links.

```
[edit]
interfaces {
  ae0 {
    aggregated-ether-options {
      link-protection;
      minimum-links 1;
      link-speed 1g;
    }
    unit 0 {
      demux-source inet {
      family inet {
        address 20.1.1.0/24;
      }
    }
    unit 1 {
      demux-source inet {
      family inet {
        address 20.1.1.1/24;
      }
    }
  }
}
```

4. Configure the IP demux interface over the aggregated Ethernet logical interface.

```
[edit]
interfaces {
  demux0 {
    unit 101 {
      demux-options {
        underlying-interface ae0.0;
      }
      family inet {
        demux-source 10.1.0.0/16;
        address 1.1.1.0/24;
      }
    }
    unit 101 {
      demux-options {
        underlying-interface ae0.1;
      }
      family inet {
        demux-source 10.1.0.1/16;
        address 1.1.1.1/24;
      }
    }
  }
}
```

Related Documentation

- Subscriber Interfaces and Demultiplexing Overview on page 759
- Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview on page 816
- Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet on page 825

Example: Configuring a Static Subscriber Interface on a VLAN Interface over Aggregated Ethernet

This example shows how you can configure a subscriber interface using a static virtual LAN (VLAN) stacked on a two-link aggregated Ethernet logical interface. In this example, the underlying aggregated Ethernet logical interface is configured for one-to-one active/backup redundancy at the DPC level, and per-subscriber static hierarchical class-of-service (CoS) is configured by applying CoS parameters at the aggregated Ethernet logical interface.

1. Define the number of aggregated Ethernet interfaces on the router.

   In this example, only one aggregated Ethernet logical interface is configured on the router.

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
```

 Copyright © 2013, Juniper Networks, Inc.

```
        }
      }
    }
```

2. Configure **ae0**, a two-link aggregated Ethernet logical interface to serve as the underlying interface for the static VLAN subscriber interface. In order to support hierarchical CoS, the physical ports must be on EQ DPCs in MX Series routers.

   In this example, the LAG bundle is configured for one-to-one active/backup link redundancy. To support link redundancy at the DPC level, the LAG bundle attaches ports from two different EQ DPCs.

```
[edit]
interfaces {
  ge-5/0/3 {
    gigether-options {
      802.3ad {
        ae0;
        primary;
      }
    }
    ge-5/1/2 {
      gigether-options {
        802.3ad {
          ae0;
          backup;
        }
      }
    }
  }
}
```

3. Configure **ae0** to serve as the underlying interface for the static VLAN interface.

```
[edit]
interfaces {
  ae0 {
    hierarchical-scheduler;
    aggregated-ether-options {
      link-protection;
      minimum-links 1;
      link-speed 1g;
      lacp {
        active;
      }
    }
  }
}
```

4. Configure static traffic-shaping and scheduling parameters.

```
[edit]
class-of-service {
  forwarding-classes { # Associate queue numbers with class names
    queue 0 be;
    queue 1 e;
    queue 2 af;
    queue 3 nc;
```

```
    }
    schedulers { # Define output queue properties
      scheduler_be {
        transmit-rate percent 30;
        buffer-size percent 30;
      }
      scheduler_ef {
        transmit-rate percent 40;
        buffer-size percent 40;
      }
      scheduler_af {
        transmit-rate percent 25;
        buffer-size percent 25;
      }
      scheduler_nc {
        transmit-rate percent 5;
        buffer-size percent 5;
      }
    }
    scheduler-maps { # Associate queues with schedulers
      smap_2 {
        forwarding-class be scheduler_be;
        forwarding-class ef scheduler_ef;
        forwarding-class-af scheduler_af;
        forwarding-class-nc scheduler_nc;
      }
    }
  }
```

5. Attach static CoS to the physical and logical interfaces of the aggregated Ethernet interface.

   In this example, three traffic control profiles are defined, but only two profiles are applied to the static VLAN subscriber interface over aggregated Ethernet:

   - The **tcp_for_ae_device_pir_500m** profile defines a shaping rate, and it is applied to both of the underlying physical interfaces (**ge-5/0/3** and **ge-5/1/2**).
   - The **tcp-for-ae_smap_video_pir_20m_delay_30m** profile defines a scheduler map, a shaping rate, and a delay buffer rate, and it is applied to one of the logical interfaces on the aggregated Ethernet bundle (**ae0.0**).

```
[edit]
class-of-service {
  traffic-control-profiles { # Configure traffic shaping and scheduling profiles
    tcp_for_ae_device_pir_500m {
      shaping-rate 20m;
    }
    tcp_for_ae_smap_video_pir_20m_delay_30m {
      scheduler-map smap_video;
      shaping-rate 20m;
      delay-buffer-rate 30m;
    }
    tcp_for_ae_smap_video_cir_50m_delay_75m {
      scheduler-map smap_video;
      guaranteed-rate 50m;
```

 Copyright © 2013, Juniper Networks, Inc.

```
            delay-buffer-rate 75m;
          }
        }
        interfaces { # Apply two traffic-control profiles to the LAG
          ae0 { # Two underlying physical interfaces on separate EQ DPCs
            output-traffic-control-profile tcp-for-ae_device_pir_500m;
            unit 0 { # One of the two logical interfaces on 'ae0'
              output-traffic-control-profile tcp-for-ae_smap_video_pir_20m_delay_30m;
            }
          }
        }
      }
```

Related Documentation

- Static and Dynamic VLAN Subscriber Interfaces over Aggregated Ethernet Overview on page 815
- Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet on page 824
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for Subscriber Access Overview on page 949

Example: Configuring IPv4 Static VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server

This example shows how to configure a static IPv4 VLAN demux interface with aggregated Ethernet as the underlying interface. DHCP Local Server configuration enables the association of subscribers to the VLAN demux interface by listing the aggregated Ethernet interface in the DHCP local server configuration.

To configure dynamic subscribers on VLAN demux interfaces:

1. Enable hierarchical scheduling and VLAN tagging on the underlying interface that you plan to use for any VLAN demux interfaces.

```
interfaces {
  ae1 {
    hierarchical-scheduler;
    vlan-tagging;
    aggregated-ether-options {
      minimum-links 1;
      lacp {
        active;
        periodic slow;
        link-protection {
          non-revertive;
        }
      }
    }
  }
}
```

2. Define the gigabit Ethernet interfaces that are part of the aggregated Ethernet interface.

```
interfaces {
  ge-5/0/0 {
    gigether-options {
      802.3ad ae1;
    }
  }
  ge-5/2/0 {
    gigether-options {
      802.3ad ae1;
    }
  }
}
```

3. Define the demux interface.

```
interfaces {
  demux0 {
    unit 102 {
      proxy-arp;
      vlan-id 103;
      demux-options {
        underlying-interface ae1;
      }
      family inet {
        unnumbered-address lo0.0 preferred-source-address 173.16.1.1;
      }
    }
  }
}
```

4. Define the loopback interface.

```
interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 192.16.1.1/32;
      }
    }
  }
}
```

5. Configure a dynamic profile for initial subscriber access.

```
dynamic-profiles {
  user-profile {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit" {
          family inet;
        }
      }
    }
    protocols {
      igmp {
        interface "$junos-interface-name" {
          version 3;
          immediate-leave;
```

 Copyright © 2013, Juniper Networks, Inc.

```
                promiscuous-mode;
            }
        }
    }
  }
}
```

6. Configure the access method used to dynamically create the subscriber interfaces.

   The following stanza specifies the aggregated Ethernet interface (**ae1.0**) for use with the dynamically created subscriber interfaces.

```
system {
  services {
    dhcp-local-server {
      group myDhcpGroup {
        authentication {
          password test;
          username-include {
            user-prefix igmp-user1;
          }
        }
        dynamic-profile user-profile;
        interface ae1.0;
      }
    }
  }
}
```

   Instead of using the aggregated Ethernet interface, you can alternatively specify the specific demux interface (**demux0.102**) as the device to use with the subscriber interfaces as follows:

```
system {
  services {
    dhcp-local-server {
      group myDhcpGroup {
        authentication {
          password test;
          username-include {
            user-prefix igmp-user1;
          }
        }
        dynamic-profile user-profile;
        interface demux0.102;
      }
    }
  }
}
```

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

Example: Separating Targeted Multicast Traffic for Demux Subscribers on Aggregated Ethernet Interfaces

This example shows how to separate targeted multicast traffic from targeted unicast traffic and send that multicast traffic to a different interface through the use of OIF maps.

- Requirements on page 846
- Overview on page 846
- Configuration on page 846
- Verification on page 851

Requirements

Before configuring this example, make sure to configure the distribution type for the interface. See "Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces" on page 828 for instructions.

Overview

In this example, targeted traffic distribution is already configured on the router. Dynamically created interfaces each carry their unicast traffic but all multicast traffic is sent to the GE-5/3/9.0 interface.

Figure 16 on page 846 shows the sample network.

Figure 16: Multicast Traffic Separation Using OIF Mapping

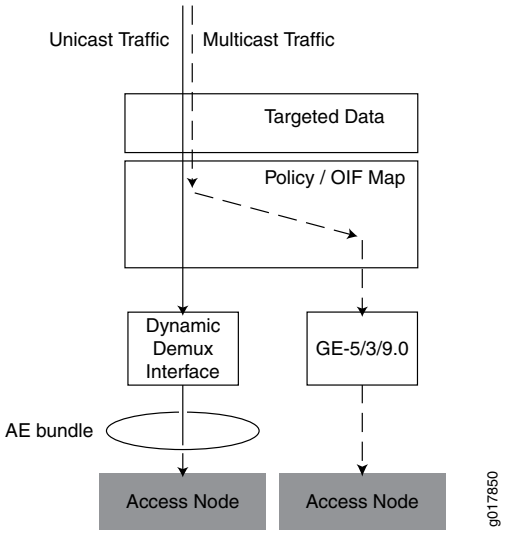

Configuration

- Configure an OIF Map Policy on page 847
- Configure a DHCP VLAN Dynamic Profile on page 848
- Configure a VLAN Demux Dynamic Profile on page 849

Copyright © 2013, Juniper Networks, Inc.

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.

```
set policy-options policy-statement OIF-v4-all term oif539 from route-filter 224 .0.0.0/4
  orlonger
set policy-options policy-statement OIF-v4-all term oif539 then map-to-interface
  ge-5/3/9.0
set policy-options policy-statement OIF-v4-all term oif539 then accept
set dynamic-profiles dhcp-vlan-prof interfaces "$junos-interface-ifd-name" unit
  "$junos-underlying-interface-unit" family inet unnumbered-address lo0.0
set dynamic-profiles dhcp-vlan-prof interfaces "$junos-interface-ifd-name" unit
  "$junos-underlying-interface-unit" family inet unnumbered-address preferred-sour
  ce-address 100.20.0.2
set dynamic-profiles demux-vlan-prof interfaces demux0 unit "$junos-interface-un it"
  vlan-id "$junos-vlan-id"
set dynamic-profiles demux-vlan-prof interfaces demux0 unit "$junos-interface-un it"
  demux-options underlying-interface "$junos-interface-ifd-name"
set dynamic-profiles demux-vlan-prof interfaces demux0 unit "$junos-interface-un it"
  targetted-distribution
set dynamic-profiles demux-vlan-prof interfaces demux0 unit "$junos-interface-un it"
  family inet unnumbered-address lo0.0
set dynamic-profiles demux-vlan-prof interfaces demux0 unit "$junos-interface-un it"
  family inet unnumbered-address preferred-source-address 100.20.0.2
set dynamic-profiles demux-vlan-prof protocols igmp interface "$junos-interface- name"
  version 2
set dynamic-profiles demux-vlan-prof protocols igmp interface "$junos-interface- name"
  promiscuous-mode
set dynamic-profiles demux-vlan-prof protocols igmp interface "$junos-interface- name"
  passive allow-receive
set dynamic-profiles demux-vlan-prof protocols igmp interface "$junos-interface- name"
  passive send-group-query
set dynamic-profiles demux-vlan-prof protocols igmp interface "$junos-interface- name"
  oif-map OIF-v4-all
```

Configure an OIF Map Policy

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the OIF map:

1. Access the router policy options:

   ```
   [edit]
   user@host#edit policy-options
   ```

2. Edit a policy statement.

   ```
   [edit policy-options]
   user@host edit policy-statement OIF-v4-all
   ```

3. Create a term for mapping incoming multicast traffic to a specific interface.

   ```
   [edit policy-options OIF-v4-all]
   ```

```
user@host edit term oif539
```

4. Define the match condition for the term. In this case, the term matches any route prefix of 224/4 or longer (all multicast traffic).

```
[edit policy-options OIF-v4-all term oif539]
user@host set from route-filter 224/4 orlonger
```

5. Define the action for the term. In this case, when a match occurs, the term accepts the traffic and maps it to interface GE-5/3/9.0.

```
[edit policy-options OIF-v4-all term oif539]
user@host set then map-to-interface ge-5/3/9.0
user@host set then accept
```

Results Confirm your configuration by issuing the **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show policy-options
policy-statement OIF-v4-all {
   term oif539 {
      from {
         route-filter 224.0.0.0/4 orlonger;
      }
      then {
         map-to-interface ge-5/3/9.0;
         accept;
      }
   }
}
```

Configure a DHCP VLAN Dynamic Profile

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a DHCP VLAN dynamic profile for client access:

1. Create a dynamic VLAN demux profile.

```
[edit]
user@host#edit dynamic-profiles dhcp-vlan-prof
```

2. Edit the dynamic profile interface.

```
[edit dynamic-profiles dhcp-vlan-prof]
user@host edit interfaces $junos-ifd-name
```

3. Edit the interface unit dynamic variable.

```
[edit dynamic-profiles demux-vlan-prof interfaces $junos-ifd-name]
user@host edit unit $junos-underlying-interface-unit
```

 Copyright © 2013, Juniper Networks, Inc.

4. Edit the interface family.

```
[edit dynamic-profiles demux-vlan-prof interfaces $junos-ifd-name unit
   $junos-underlying-interface-unit]
user@host edit family inet
```

5. Define the loopback address.

```
[edit dynamic-profiles demux-vlan-prof interfaces $junos-ifd-name unit
   $junos-underlying-interface-unit ]
user@host set unnumbered-address lo0.0 preferred-source-address 100.20.0.2
```

Results Confirm your configuration by issuing the **show dynamic-profiles** command. If the output for the dhcp-vlan-prof dynamic profile does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show dynamic-profiles
dhcp-vlan-prof {
   interfaces {
      "$junos-interface-ifd-name" {
         unit "$junos-underlying-interface-unit" {
            family inet {
               unnumbered-address lo0.0 preferred-source-address 100.20.0.2;
            }
         }
      }
   }
}
```

Configure a VLAN Demux Dynamic Profile

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the OIF map:

1. Create a dynamic VLAN demux profile.

```
[edit]
user@host#edit dynamic-profiles demux-vlan-prof
```

2. Edit the dynamic profile demux0 interface.

```
[edit dynamic-profiles demux-vlan-prof]
user@host edit interfaces demux0
```

3. Edit the interface unit dynamic variable.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0]
user@host edit unit $junos-interface-unit
```

4. Specify the VLAN ID dynamic variable.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit
   “$junos-interface-unit”]
```

```
user@host set vlan-id $junos-vlan-id
```

5. Access the demux options.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit
  "$junos-interface-unit"]
user@host edit demux-options
```

6. Define the demux underlying interface.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit
  "$junos-interface-unit" demux-options]
user@host set underlying-interface $junos-interface-ifd-name
```

7. Specify that dynamically created VLANs are using targeted distribution.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit
  "$junos-interface-unit"]
user@host set targeted-distribution
```

8. Edit the interface family.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit
  "$junos-interface-unit"]
user@host edit family inet
```

9. Define the loopback address.

```
[edit dynamic-profiles demux-vlan-prof interfaces demux0 unit
  "$junos-interface-unit" family inet]
user@host set unnumbered-address lo0.0 preferred-source-address 100.20.0.2
```

10. Edit the dynamic profile IGMP protocol.

```
[edit dynamic-profiles demux-vlan-prof]
user@host edit protocols igmp
```

11. Enable IGMP on dynamically created interfaces.

```
[edit dynamic-profiles demux-vlan-prof protocols igmp]
user@host edit interface $junos-interface-name
```

12. Specify the IGMP version that you want dynamically created interfaces to use.

```
[edit dynamic-profiles demux-vlan-prof protocols igmp interface
  $junos-interface-name]
user@host set version 2
```

13. Specify the OIF map that you want dynamically created IGMP interfaces to use.

```
[edit dynamic-profiles demux-vlan-prof protocols igmp interface
  $junos-interface-name]
user@host set oif-map OIF-v4-all
```

14. Specify that IGMP selectively sends and receives control traffic such as IGMP reports, queries, and leaves.

```
[edit dynamic-profiles demux-vlan-prof protocols igmp interface
  $junos-interface-name]
user@host set passive allow-receive send-group-query
```

 Copyright © 2013, Juniper Networks, Inc.

15. Specify that the interface accepts IGMP reports from hosts on any subnetwork.

```
[edit dynamic-profiles demux-vlan-prof protocols igmp interface
  $junos-interface-name]
user@host set promiscuous-mode
```

Results Confirm your configuration by issuing the **show dynamic-profiles** commands. If the output for the dhcp-vlan-prof dynamic profile does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show dynamic-profiles
demux-vlan-prof {
  interfaces {
    demux0 {
      unit "$junos-interface-unit" {
        vlan-id "$junos-vlan-id";
        demux-options {
          underlying-interface "$junos-interface-ifd-name";
        }
        targetted-distribution;
        family inet {
          unnumbered-address lo0.0 preferred-source-address 100.20.0.2;
        }
      }
    }
  }
  protocols {
    igmp {
      interface "$junos-interface-name" {
        version 2;
        promiscuous-mode;
        passive allow-receive send-group-query;
        oif-map OIF-v4-all;
      }
    }
  }
}
...
```

Verification

Confirm that the configuration is working properly.

- Locate the Multicast Group Member on page 851
- Ensure the Targeting Aggregated Ethernet Interface for the Subscriber is Functional on page 852
- View the Packets for the Targeted Interface on page 853

Locate the Multicast Group Member

Purpose Locate the dynamic interface and ensure that it is associated with the appropriate IGMP group.

Action `user@host>`**show igmp group**

```
Interface: demux0.1073741824, Groups: 1
    Group: 225.0.0.1
        Source: 0.0.0.0
        Last reported by: 100.20.0.10
        Timeout:      52 Type: Dynamic
Interface: local, Groups: 2
    Group: 224.0.0.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

Meaning The first **Interface** field shows the dynamically created demux interface, **demux0.1073741824**, and the Group field immediately below the first Interface field shows the group, **225.0.0.1**, to which the subscriber belongs.

Ensure the Targeting Aggregated Ethernet Interface for the Subscriber is Functional

Purpose Use the dynamic subscriber interface value to ensure that the targeting aggregated interface is functional.

 Copyright © 2013, Juniper Networks, Inc.

Action `user@host>` **show interfaces demux0.1073741824 extensive**

```
Logical interface demux0.1073741824 (Index 810) (SNMP ifIndex 1613)
(Generation 170)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
  Demux:
    Underlying interface: ae0 (Index 708)
  Link:
    ge-1/0/0
    ge-5/3/7
  Targeting summary:
    ge-1/0/0, backup, Physical link is Up
    ge-5/3/7, primary, Physical link is Up
  Traffic statistics:
   Input  bytes  :                  862
   Output bytes  :                 3160
   Input  packets:                    3
   Output packets:                   30
  Local statistics:
   Input  bytes  :                  862
   Output bytes  :                 3160
   Input  packets:                    3
   Output packets:                   30
  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: 212, Route table: 0
    Flags: Sendbcast-pkt-to-re, Unnumbered
    Donor interface: lo0.0 (Index 802)
    Preferred source address: 100.20.0.2
```

Meaning The **Targeting summary** field shows that the primary interface, **ge-5/3/7**, is up.

View the Packets for the Targeted Interface

Purpose Verify that packet traffic sent to targeted interface GE-5/3/9 consists only of multicast packets.

Action

```
user@host>show interfaces ge-5/3/9 extensive
Physical interface: ge-5/3/9, Enabled, Physical link is Up
  Interface index: 704, SNMP ifIndex: 1605, Generation: 197
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:21:59:ab:85:2a, Hardware address: 00:21:59:ab:85:2a
  Last flapped   : 2012-09-26 17:32:24 EDT (6d 20:44 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :             97857650                 1320 bps
   Output bytes  :                    0                    0 bps
   Input  packets:               889615                    1 pps
   Output packets:                    0               889620 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  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: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0        113871616
    Total packets                            0           889620
    Unicast packets                          0                0
    Broadcast packets                        0                0
```


Copyright © 2013, Juniper Networks, Inc.

```
  Multicast packets                        0           889620
  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
  Total errors                             0                0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                  889620
  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, Remote fault: OK
  Local resolution:
      Flow control: None, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue            Bandwidth               Buffer Priority   Limit

                         %            bps     %            usec
  0 best-effort         95      950000000    95               0      low    none

  3 network-control      5       50000000     5               0      low    none

Interface transmit statistics: Disabled

Logical interface ge-5/3/9.0 (Index 818) (SNMP ifIndex 1597) (Generation 149)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :             97857650
   Input  packets:                    0
   Output packets:               889620
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Transit statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :             97857650                 1320 bps
   Input  packets:                    0                    0 pps
   Output packets:               889615                    1 pps
  Protocol aenet, AE bundle: ae4.0, Generation: 180, Route table: 0
```

Meaning The MAC statistics **Unicast packet** field shows that the interface is not transmitting any

unicast packet traffic and the **Multicast packet** field shows that the total number of packets being transmitted from the interface are multicast packets.

Related Documentation

- Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces on page 828

Example: Configuring IPv4 Dynamic VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server

This example shows how to configure the dynamic creation of IPv4 VLAN demux interfaces with aggregated Ethernet as the underlying interface. DHCP Local Server configuration enables the association of subscribers to the VLAN demux interface by listing the aggregated Ethernet interface in the DHCP local server configuration.

NOTE: VLAN demux subscriber interfaces over aggregated Ethernet physical interfaces are supported only for MX Series routers that have only MPCs installed. If the router has other cards in addition to MPCs, the CLI accepts the configuration but errors are reported when the subscriber interfaces are brought up.

To configure dynamic subscribers on dynamic VLAN demux interfaces:

1. Enable VLAN tagging and VLAN auto-configuration on the underlying aggregated Ethernet interface that you plan to use for dynamically created VLAN demux interfaces.

```
interfaces {
  ae1 {
    vlan-tagging;
    auto-configure {
      vlan-ranges {
        dynamic-profile auto-vlanDemux-profile {
          accept inet;
          ranges {
            any;
          }
        }
      }
    }
    aggregated-ether-options {
      minimum-links 1;
      lacp {
        active;
        periodic slow;
        link-protection {
          non-revertive;
        }
      }
    }
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

2. Define the gigabit Ethernet interfaces that are part of the aggregated Ethernet interface.

```
interfaces {
  ge-5/0/0 {
    gigether-options {
      802.3ad ae1;
    }
  }
  ge-5/2/0 {
    gigether-options {
      802.3ad ae1;
    }
  }
}
```

3. Define the loopback interface.

```
interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 192.16.1.1/32;
      }
    }
  }
}
```

4. Configure a dynamic profile for subscriber access.

```
dynamic-profiles {
  user-profile {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit" {
          family inet;
        }
      }
    }
  }
}
```

5. Configure a dynamic profile for VLAN demux interface creation.

```
dynamic-profiles {
  auto-vlanDemux-profile {
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          vlan-id "$junos-vlan-id";
          demux-options {
            underlying-interface "$junos-interface-ifd-name";
          }
          family inet {
            filter {
              input rate_limit;
              output rate_limit;
            }
            unnumbered-address lo0.0 preferred-source-address 192.16.1.1;
```

```
            }
          }
        }
      }
    }
  }
```

6. Configure the access method used to dynamically create the subscriber interfaces. The following stanza specifies the aggregated Ethernet interface (**ae1.0**) for use with the dynamically created subscriber interfaces.

```
system {
  services {
    dhcp-local-server {
      group myDhcpGroup {
        authentication {
          password test;
          username-include {
            user-prefix igmp-user1;
          }
        }
        dynamic-profile user-profile;
        interface ae1.0;
      }
    }
  }
}
```

Instead of using the aggregated Ethernet interface, you can alternatively specify **demux0** as the device to use with the subscriber interfaces as follows:

NOTE: Because the demux interfaces and unit values are created dynamically, the unit number is not specified for the demux0 interface.

```
system {
  services {
    dhcp-local-server {
      group myDhcpGroup {
        authentication {
          password test;
          username-include {
            user-prefix igmp-user1;
          }
        }
        dynamic-profile user-profile;
        interface demux0;
      }
    }
  }
}
```

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles on page 772

Copyright © 2013, Juniper Networks, Inc.

- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

Example: Configuring IPv6 Dynamic VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server

This example shows how to configure the dynamic creation of IPv6 VLAN demux interfaces with aggregated Ethernet as the underlying interface. DHCP Local Server configuration enables the association of subscribers to the VLAN demux interface by listing the aggregated Ethernet interface in the DHCP local server configuration.

NOTE: VLAN demux subscriber interfaces over aggregated Ethernet physical interfaces are supported only for MX Series routers that have only MPCs installed. If the router has other cards in addition to MPCs, the CLI accepts the configuration but errors are reported when the subscriber interfaces are brought up.

To configure dynamic subscribers on dynamic VLAN demux interfaces:

1. Enable VLAN tagging and VLAN auto-configuration on the underlying aggregated Ethernet interface that you plan to use for dynamically created VLAN demux interfaces.

```
interfaces {
  ae1 {
    vlan-tagging;
    auto-configure {
      vlan-ranges {
        dynamic-profile auto-vlanDemux-profile {
          accept inet6;
          ranges {
            any;
          }
        }
      }
    }
    aggregated-ether-options {
      minimum-links 1;
      lacp {
        active;
        periodic slow;
        link-protection {
          non-revertive;
        }
      }
    }
  }
}
```

2. Define the gigabit Ethernet interfaces that are part of the aggregated Ethernet interface.

```
interfaces {
  ge-5/0/0 {
```

```
        gigether-options {
            802.3ad ae1;
        }
    }
    ge-5/2/0 {
        gigether-options {
            802.3ad ae1;
        }
    }
}
```

3. Define the loopback interface.

```
interfaces {
    lo0 {
        unit 0 {
            family inet6 {
                address 2009:174:1:1::1/128;
            }
        }
    }
}
```

4. Configure a dynamic profile for subscriber access.

```
dynamic-profiles {
    user-profile {
        interfaces {
            "$junos-interface-ifd-name" {
                unit "$junos-underlying-interface-unit" {
                    family inet6;
                }
            }
        }
    }
}
```

5. Configure a dynamic profile for VLAN demux interface creation.

```
dynamic-profiles {
    auto-vlanDemux-profile {
        interfaces {
            demux0 {
                unit "$junos-interface-unit" {
                    vlan-id "$junos-vlan-id";
                    demux-options {
                        underlying-interface "$junos-interface-ifd-name";
                    }
                    family inet6 {
                        filter {
                            input v6_rate_limit;
                            output v6_rate_limit;
                        }
                        unnumbered-address lo0.0 preferred-source-address 2009:174:1:1::1;
                    }
                }
            }
        }
```

 Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
```

6. Configure the access method used to dynamically create the subscriber interfaces. The following stanza specifies the aggregated Ethernet interface (**ae1.0**) for use with the dynamically created subscriber interfaces.

```
system {
  services {
    dhcp-local-server {
      dhcpv6 {
        group myV6DhcpGroup {
          authentication {
            password test;
            username-include {
              user-prefix igmp-user1;
            }
          }
          dynamic-profile user-profile;
          interface ae1.0;
        }
      }
    }
  }
}
```

Instead of using the aggregated Ethernet interface, you can alternatively specify **demux0** as the device to use with the subscriber interfaces as follows:

NOTE: Because the demux interfaces and unit values are created dynamically, the unit number is not specified for the demux0 interface.

```
system {
  services {
    dhcp-local-server {
      dhcpv6 {
        group myV6DhcpGroup {
          authentication {
            password test;
            username-include {
              user-prefix igmp-user1;
            }
          }
          dynamic-profile user-profile;
          interface demux0;
        }
      }
    }
  }
}
```

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles on page 772

• Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

Example: Configuring IPv4 Dynamic Stacked VLAN Demux Interfaces over an Aggregated Ethernet Underlying Interface with DHCP Local Server

This example shows how to configure the dynamic creation of IPv4 stacked VLAN demux interfaces with aggregated Ethernet as the underlying interface. DHCP Local Server configuration enables the association of subscribers to the VLAN demux interface by listing the aggregated Ethernet interface in the DHCP local server configuration.

NOTE: VLAN demux subscriber interfaces over aggregated Ethernet physical interfaces are supported only for MX Series routers that have only MPCs installed. If the router has other cards in addition to MPCs, the CLI accepts the configuration but errors are reported when the subscriber interfaces are brought up.

To configure dynamic subscribers on dynamic VLAN demux interfaces:

1. Enable VLAN tagging and VLAN auto-configuration on the underlying aggregated Ethernet interface that you plan to use for dynamically created VLAN demux interfaces.

```
interfaces {
  ae1 {
    flexible-vlan-tagging;
    auto-configure {
      stacked-vlan-ranges {
        dynamic-profile auto-vlanDemux-profile {
          accept inet;
          ranges {
            any;
          }
        }
      }
    }
    aggregated-ether-options {
      minimum-links 1;
      lacp {
        active;
        periodic slow;
        link-protection {
          non-revertive;
        }
      }
    }
  }
}
```

2. Define the gigabit Ethernet interfaces that are part of the aggregated Ethernet interface.

```
interfaces {
  ge-5/0/0 {
```

 Copyright © 2013, Juniper Networks, Inc.

```
        gigether-options {
          802.3ad ae1;
        }
      }
      ge-5/2/0 {
        gigether-options {
          802.3ad ae1;
        }
      }
    }
```

3. Define the loopback interface.

```
interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 192.16.1.1/32;
      }
    }
  }
}
```

4. Configure a dynamic profile for subscriber access.

```
dynamic-profiles {
  user-profile {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit" {
          family inet;
        }
      }
    }
  }
}
```

5. Configure a dynamic profile for VLAN demux interface creation.

```
dynamic-profiles {
  auto-vlanDemux-profile {
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          vlan-tags outer "$junos-stacked-vlan-id" inner "$junos-vlan-id";
          demux-options {
            underlying-interface "$junos-interface-ifd-name";
          }
          family inet {
            filter {
              input rate_limit;
              output rate_limit;
            }
            unnumbered-address lo0.0 preferred-source-address 192.16.1.1;
          }
        }
      }
    }
```

```
    }
  }
```

6. Configure the access method used to dynamically create the subscriber interfaces. The following stanza specifies the aggregated Ethernet interface (**ae1.0**) for use with the dynamically created subscriber interfaces.

```
system {
  services {
    dhcp-local-server {
      group myDhcpGroup {
        authentication {
          password test;
          username-include {
            user-prefix igmp-user1;
          }
        }
        dynamic-profile user-profile;
        interface ae1.0;
      }
    }
  }
}
```

Instead of using the aggregated Ethernet interface, you can alternatively specify **demux0** as the device to use with the subscriber interfaces as follows:

NOTE: Because the demux interfaces and unit values are created dynamically, the unit number is not specified for the demux0 interface.

```
system {
  services {
    dhcp-local-server {
      group myDhcpGroup {
        authentication {
          password test;
          username-include {
            user-prefix igmp-user1;
          }
        }
        dynamic-profile user-profile;
        interface demux0;
      }
    }
  }
}
```

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles on page 772
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring a Static PPPoE Subscriber Interface on a Static Underlying VLAN Demux Interface over Aggregated Ethernet

This example shows how you can configure static PPPoE subscriber interfaces over aggregated Ethernet bundles to provide subscriber link redundancy.

- Requirements on page 865
- Overview on page 865
- Configuration on page 865
- Verification on page 868

Requirements

PPPoE over VLAN demux interfaces over aggregated Ethernet requires the following hardware and software:

- MX Series 3D Universal Edge Routers
- MPCs
- Junos OS Release 11.2 or later

No special configuration beyond device initialization is required before you can configure this feature.

Overview

Aggregated Ethernet bundles enable link redundancy between the router and networking devices connected by Ethernet links. This example describes how to configure link redundancy for static PPPoE subscribers over aggregated Ethernet interface with an intermediate static VLAN demux interface. Sample tasks include configuring a two-member aggregated Ethernet bundle on **ae0**, configuring a static VLAN demux interface, **demux0.100**, that underlies the PPPoE subscriber interface, **pp0.100**, and configuring the PPPoE subscriber interface including characteristics of the PPPoE family.

This example does not show all possible configuration choices.

Configuration

CLI Quick Configuration

To quickly configure link redundancy for static PPPoE subscribers over a static VLAN demux interface over aggregated Ethernet, 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]
set chassis aggregated-devices ethernet device-count 1
set interfaces ge-5/0/3 gigether-options 802.3ad ae0
set interfaces ge-5/0/3 gigether-options 802.3ad primary
set interfaces ge-5/1/2 gigether-options 802.3ad ae0
set interfaces ge-5/1/2 gigether-options 802.3ad backup
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 aggregated-ether-options link-protection
```

```
edit interfaces demux0 unit 100
set vlan-id 100
set demux-options underlying-interface ae0
set family pppoe access-concentrator pppoe-server-1
set family pppoe duplicate-protection
set family pppoe max-sessions 16000
top
edit interfaces pp0 unit 100
set pppoe-options underlying-interface demux0.100
set pppoe-options server
set family inet unnumbered-address lo0.0
top
```

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 link redundancy for static PPPoE subscribers over a static VLAN demux interface over aggregated Ethernet:

1. Define the number of aggregated Ethernet devices on the router.

   ```
   [edit chassis]
   user@host# set aggregated-devices ethernet device-count 1
   ```

2. Configure a two-link aggregated Ethernet logical interface to serve as the underlying interface for the static VLAN demux subscriber interface. In this example, the LAG bundle is configured for one-to-one active/backup link redundancy. To support link redundancy at the MPC level, the LAG bundle attaches to ports from two different MPCs.

   ```
   [edit interfaces]
   user@host# set ge-5/0/3 gigether-options 802.3ad ae0
   user@host# set ge-5/0/3 gigether-options 802.3ad primary
   user@host# set ge-5/1/2 gigether-options 802.3ad ae0
   user@host# set ge-5/1/2 gigether-options 802.3ad backup
   ```

3. Enable link protection on the aggregated Ethernet logical interface and configure support for single and dual (stacked) VLAN tags.

   ```
   [edit interfaces]
   user@host# set ae0 aggregated-ether-options link-protection
   user@host# set ae0 flexible-vlan-tagging
   ```

4. Configure the VLAN demux interface over the aggregated Ethernet logical interface.

   ```
   [edit interfaces]
   user@host# set demux0 unit 100 vlan-id 100
   user@host# set demux0 unit 100 demux-options underlying-interface ae0
   ```

5. Configure the PPPoE family attributes on the VLAN demux interface.

   ```
   [edit interfaces]
   user@host# set demxu0 unit 100 family pppoe access-concentrator pppoe-server-1
   user@host# set demux0 unit 100 family pppoe duplicate-protection
   user@host# set demux0 unit 100 family pppoe max-sessions 16000
   ```


Copyright © 2013, Juniper Networks, Inc.

6. Configure the VLAN demux interface as the underlying interface on which the PPPoE logical interface is created.

```
[edit interfaces]
user@host# set pp0 unit 100 pppoe-options underlying-interface demux0.100
user@host# set pp0 unit 100 pppoe-options server
user@host# set pp0 unit 100 family inet unnumbered-address lo0.0
```

Results From configuration mode, confirm the aggregated device configuration by entering the **show chassis** command. Confirm the interface configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show chassis
aggregated-devices {
  ethernet {
    device-count 1;
  }
}
```

```
[edit]
user@host# show interfaces
  ge-5/0/3 {
    gigether-options {
      802.3ad {
        ae0;
        primary;
      }
    }
  }
  ge-5/1/2 {
    gigether-options {
      802.3ad {
        ae0;
        backup;
      }
    }
  }
  ae0 {
    flexible-vlan-tagging;
    aggregated-ether-options {
      link-protection;
    }
  }
  demux0 {
    unit 100 {
      vlan-id 100;
      demux-options {
        underlying-interface ae0;
      }
      family pppoe {
        access-concentrator pppoe-server-1;
        duplicate-protection;
        max-sessions 16000;
```

```
        }
      }
    }
    pp0 {
      unit 100 {
        pppoe-options {
          underlying-interface demux0.100;
          server;
        }
        family inet {
          unnumbered-address lo0.0;
        }
      }
    }
```

If 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 868
- Verifying the demux0 Interface Configuration on page 869
- Verifying the pp0 Interface Configuration on page 869

Verifying the Aggregated Ethernet Interface Configuration

Purpose Verify that the interface values match your configuration, the link is up, and traffic is flowing.

Action From operational mode, enter the **show interfaces redundancy** command.

```
user@host> show interfaces redundancy
Interface  State           Last change  Primary      Secondary    Current status
ae0        On primary                   ge-5/0/3     ge-5/1/2     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: 00:1f:12:b8:ef:c0, Hardware address: 00:1f:12:b8:ef:c0
  Last flapped   : 2011-03-11 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
```

 Copyright © 2013, Juniper Networks, Inc.

```
    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 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.

```
user@host> show interfaces demux0.100
Logical interface demux0.100 (Index 76) (SNMP ifIndex 61160)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
    Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 199)
    Link:
      ge-5/0/3
      ge-5/1/2
    Input packets : 2
    Output packets: 18575
    Protocol pppoe
      Dynamic Profile: none,
      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: 00:90:1a:00:18:34,
      Underlying interface: demux0.100 (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: 45.63.24.1
```

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.

Related Documentation

- Subscriber Interfaces and Demultiplexing Overview on page 759
- Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview on page 816
- Configuring Static Subscriber Interfaces Using VLAN Demux Interfaces on page 767
- Configuring the PPPoE Family for an Underlying Interface on page 831

Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Underlying VLAN Demux Interface over Aggregated Ethernet

This example shows how you can configure dynamic PPPoE subscriber interfaces over aggregated Ethernet bundles to provide subscriber link redundancy.

- Requirements on page 870
- Overview on page 871
- Configuration on page 871
- Verification on page 874

Requirements

PPPoE over VLAN demux interfaces over aggregated Ethernet requires the following hardware and software:

- MX Series 3D Universal Edge Routers
- MPCs
- Junos OS Release 11.2 or later

No special configuration beyond device initialization is required before you can configure this feature.

 Copyright © 2013, Juniper Networks, Inc.

Overview

Aggregated Ethernet bundles enable link redundancy between the router and networking devices connected by Ethernet links. This example describes how to configure link redundancy for dynamic PPPoE subscribers over aggregated Ethernet interface, **ae0**, with an intermediate static VLAN demux interface, **demux0.100**. Sample tasks include configuring a two-member aggregated Ethernet bundle, configuring a static VLAN demux interface that underlies the PPPoE subscriber interface, and configuring the dynamic profile that establishes the dynamic PPPoE subscriber interfaces.

The dynamic PPPoE profile (**pppoe-profile**) creates the PPPoE subscriber interface. It also configures the router to act as a PPPoE server and enables the local address to be derived from the specified address without assigning an explicit IP address to the interface. The **pppoe-profile** dynamic profile is assigned to the static, intermediate VLAN demux interface (**demux0.100**), which is configured with the PPPoE family (**family pppoe**) attributes. This dynamic profile includes the following predefined variables:

- **$junos-interface-unit**—Represents the logical unit number of the dynamic PPPoE logical interface. This predefined variable is dynamically replaced with the unit number supplied by the router when the subscriber logs in.
- **$junos-underlying-interface**—Represents the name of the underlying Ethernet interface. This predefined variable is dynamically replaced with the interface name supplied by the router when the subscriber logs in.

This example does not show all possible configuration choices.

Configuration

CLI Quick Configuration

To quickly configure link redundancy for dynamic PPPoE subscribers over a static VLAN demux interface over aggregated Ethernet, 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]
set chassis aggregated-devices ethernet device-count 1
set interfaces ge-5/0/3 gigether-options 802.3ad ae0
set interfaces ge-5/0/3 gigether-options 802.3ad primary
set interfaces ge-5/1/2 gigether-options 802.3ad ae0
set interfaces ge-5/1/2 gigether-options 802.3ad backup
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 aggregated-ether-options link-protection
set interfaces demux0 unit 100 vlan-id 100
set interfaces demux0 unit 100 demux-options underlying-interface ae0
set interfaces demux0 unit 100 family pppoe access-concentrator pppoe-server-1
set interfaces demux0 unit 100 family pppoe duplicate-protection
set interfaces demux0 unit 100 family pppoe dynamic-profile pppoe-profile
edit dynamic-profiles pppoe-profile
edit interfaces pp0 unit $junos-interface-unit
set pppoe-options underlying-interface $junos-underlying-interface
set pppoe-options server
set family inet unnumbered-address lo0.0
top
```

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 link redundancy for dynamic PPPoE subscribers over a static VLAN demux interface over aggregated Ethernet:

1. Define the number of aggregated Ethernet devices on the router.

```
[edit chassis]
user@host# set aggregated-devices ethernet device-count 1
```

2. Configure a two-link aggregated Ethernet logical interface to serve as the underlying interface for the static VLAN demux subscriber interface. In this example, the LAG bundle is configured for one-to-one active/backup link redundancy. To support link redundancy at the MPC level, the LAG bundle attaches to ports from two different MPCs.

```
[edit interfaces]
user@host# set ge-5/0/3 gigether-options 802.3ad ae0
user@host# set ge-5/0/3 gigether-options 802.3ad primary
user@host# set ge-5/1/2 gigether-options 802.3ad ae0
user@host# set ge-5/1/2 gigether-options 802.3ad backup
```

3. Enable link protection on the aggregated Ethernet logical interface and configure support for single and dual (stacked) VLAN tags.

```
[edit interfaces]
user@host# set ae0 aggregated-ether-options link-protection
user@host# set ae0 flexible-vlan-tagging
```

4. Configure the VLAN demux interface over the aggregated Ethernet logical interface.

```
[edit interfaces]
user@host# set demux0 unit 100 vlan-id 100
user@host# set demux0 unit 100 demux-options underlying-interface ae0
```

5. Configure the PPPoE family attributes on the VLAN demux interface, including the dynamic profile.

```
[edit interfaces]
user@host# set demux0 unit 100 family pppoe access-concentrator pppoe-server-1
user@host# set demux0 unit 100 family pppoe duplicate-protection
user@host# set demux0 unit 100 family pppoe dynamic-profile pppoe-profile
```

6. Configure the dynamic profile that creates the PPPoE subscriber interfaces.

```
[edit dynamic-profiles pppoe-profile]
user@host# edit interfaces pp0 unit $junos-interface-unit
[edit dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options underlying-interface $junos-underlying-interface
user@host# set pppoe-options server
user@host# set family inet unnumbered-address lo0.0
```

Results

From configuration mode, confirm the aggregated device configuration by entering the **show chassis** command. Confirm the interface configuration by entering the **show interfaces** command. Confirm the dynamic profile configuration by entering the **show**

Copyright © 2013, Juniper Networks, Inc.

dynamic-profiles command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show chassis
aggregated-devices {
   ethernet {
      device-count 1;
   }
}
```

```
[edit]
user@host# show interfaces
   ge-5/0/3 {
      gigether-options {
         802.3ad {
            ae0;
            primary;
         }
      }
   }
   ge-5/1/2 {
      gigether-options {
         802.3ad {
            ae0;
            backup;
         }
      }
   }
   ae0 {
      flexible-vlan-tagging;
      aggregated-ether-options {
         link-protection;
      }
   }
   demux0 {
      unit 100 {
         vlan-id 100;
         demux-options {
            underlying-interface ae0;
         }
         family pppoe {
            access-concentrator pppoe-server-1
            duplicate-protection;
            dynamic-profile pppoe-profile;
         }
      }
   }
```

```
[edit]
user@host# show dynamic-profiles
pppoe-profile {
   interfaces {
      pp0 {
         unit $junos-interface-unit {
            pppoe-options {
```

```
                underlying-interface $junos-underlying-interface;
                server;
            }
            family inet {
                unnumbered-address lo0.0;
            }
        }
    }
  }
}
```

If 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 874
- Verifying the demux0 Interface Configuration on page 875

Verifying the Aggregated Ethernet Interface Configuration

Purpose Verify that the interface values match your configuration, the link is up, and traffic is flowing.

Action From operational mode, enter the **show interfaces redundancy** command.

```
user@host> show interfaces redundancy
Interface  State           Last change  Primary      Secondary    Current status
ae0        On primary                   ge-5/0/3     ge-5/1/2     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: 00:1f:12:b8:ef:c0, Hardware address: 00:1f:12:b8:ef:c0
  Last flapped   : 2011-03-11 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
```


Copyright © 2013, Juniper Networks, Inc.

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 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.

```
user@host> show interfaces demux0.100
Logical interface demux0.100 (Index 76) (SNMP ifIndex 61160)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
    Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 199)
    Link:
      ge-5/0/3
      ge-5/1/2
    Input packets : 2
    Output packets: 18575
    Protocol pppoe
      Dynamic Profile: pppoe-profile,
      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. The output also provides information about PPPoE negotiation on a per-VLAN basis.

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.

Related Documentation

- Subscriber Interfaces and Demultiplexing Overview on page 759
- Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview on page 816
- Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles on page 772
- Configuring the PPPoE Family for an Underlying Interface on page 831
- Configuring a Basic PPPoE Dynamic Profile on page 900

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring a Dynamic PPPoE Subscriber Interface on a Dynamic Underlying VLAN Demux Interface over Aggregated Ethernet

This example shows how you can configure dynamic PPPoE subscriber interfaces over aggregated Ethernet bundles to provide subscriber link redundancy.

- Requirements on page 876
- Overview on page 876
- Configuration on page 877
- Verification on page 881

Requirements

PPPoE over VLAN demux interfaces over aggregated Ethernet requires the following hardware and software:

- MX Series 3D Universal Edge Routers
- MPCs
- Junos OS Release 11.2 or later

No special configuration beyond device initialization is required before you can configure this feature.

Overview

Aggregated Ethernet bundles enable link redundancy between the router and networking devices connected by Ethernet links. This example describes how to configure link redundancy for dynamic PPPoE subscribers over aggregated Ethernet with an intermediate dynamic VLAN demux interface. Sample tasks include configuring a two-member aggregated Ethernet bundle, configuring dynamic profiles that establish the dynamic VLAN demux interface that underlies the PPPoE subscriber interface, and configuring the dynamic profile that establishes the dynamic PPPoE subscriber interfaces.

In this example, two different dynamic profiles are configured to instantiate either VLAN (**vlan-profile**) or S-VLAN (**svlan-profile**) demux interfaces. These profiles define PPPoE family options and include the dynamic PPPoE profile (**pppoe-profile**) that creates the PPPoE subscriber interface. Junos OS predefined variables are used in each profile to represent the interfaces and VLAN identifiers that are dynamically created. These dynamic profiles include the following predefined variables:

- **$junos-interface-unit**—Represents the logical unit number of the dynamic VLAN demux interface. This predefined variable is dynamically replaced with the unit number supplied by the router when the subscriber logs in.
- **$junos-interface-ifd-name**—Represents the underlying logical interface on which the PPPoE subscriber interface is created. This predefined variable is dynamically replaced with the name of the underlying interface supplied by the router when the subscriber logs in.

Copyright © 2013, Juniper Networks, Inc.

- **$junos-vlan-id**—Represents the VLAN identifier. This predefined variable is dynamically replaced with a VLAN ID when the subscriber logs in. The VLAN ID is allocated within the VLAN range specified in the aggregated Ethernet configuration. In the case of the S-VLAN demux, **$junos-vlan-id** represents the inner VLAN identifier.
- **$junos-stacked-vlan-id**—Represents the outer VLAN identifier for the stacked VLAN. This predefined variable is dynamically replaced with a VLAN ID when the subscriber logs in. The VLAN ID is allocated within the VLAN range specified in the aggregated Ethernet configuration. This variable is not used for the VLAN demux configuration.

The dynamic PPPoE profile (**pppoe-profile**) creates the PPPoE subscriber interface. It also configures the router to act as a PPPoE server and enables the local address to be derived from the specified address without assigning an explicit IP address to the interface. The **pppoe-profile** dynamic profile is assigned to the dynamic, intermediate VLAN and S-VLAN demux interfaces. This dynamic profile includes the following predefined variables:

- **$junos-interface-unit**—Represents the logical unit number of the dynamic PPPoE logical interface. This predefined variable is dynamically replaced with the unit number supplied by the router when the subscriber logs in.
- **$junos-underlying-interface**—Represents the name of the underlying Ethernet interface. This predefined variable is dynamically replaced with the interface name supplied by the router when the subscriber logs in.

This example does not show all possible configuration choices.

Configuration

CLI Quick Configuration

To quickly configure link redundancy for dynamic PPPoE subscribers over a dynamic VLAN demux interface over aggregated Ethernet, 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]
set chassis aggregated-devices ethernet device-count 1
set interfaces ge-5/0/3 gigether-options 802.3ad ae0
set interfaces ge-5/0/3 gigether-options 802.3ad primary
set interfaces ge-5/1/2 gigether-options 802.3ad ae0
set interfaces ge-5/1/2 gigether-options 802.3ad backup
edit interfaces ae0
set flexible-vlan-tagging
set aggregated-ether-options link-protection
edit auto-configure
set vlan-ranges dynamic-profile vlan-profile accept pppoe
set vlan-ranges dynamic-profile vlan-profile ranges 1-4094
set stacked-vlan-ranges dynamic-profile svlan-profile accept pppoe
set stacked-vlan-ranges dynamic-profile svlan-profile ranges 1-4094,1-4094
top
edit dynamic-profiles pppoe-profile
edit interfaces pp0 unit $junos-interface-unit
set pppoe-options underlying-interface $junos-underlying-interface
set pppoe-options server
set family inet unnumbered-address lo0.0
```

```
top
edit dynamic-profiles vlan-profile interfaces demux0
edit unit $junos-interface-unit
set vlan-id $junos-vlan-id
set demux-options underlying-interface $junos-interface-ifd-name
set family pppoe access-concentrator pppoe-server-1
set family pppoe duplicate-protection
set family pppoe dynamic-profile pppoe-profile
top
edit dynamic-profiles svlan-profile interfaces demux0
edit unit $junos-interface-unit
set vlan-tags outer $junos-stacked-vlan-id
set vlan-tags inner $junos-vlan-id
set demux-options underlying-interface $junos-interface-ifd-name
set family pppoe access-concentrator pppoe-server-1
set family pppoe duplicate-protection
set family pppoe dynamic-profile pppoe-profile
top
```

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 link redundancy for dynamic PPPoE subscribers over a dynamic VLAN demux interface over aggregated Ethernet:

1. Define the number of aggregated Ethernet devices on the router.

   ```
   [edit chassis]
   user@host# set aggregated-devices ethernet device-count 1
   ```

2. Configure a two-link aggregated Ethernet logical interface to serve as the underlying interface for the dynamic VLAN demux subscriber interface. In this example, the LAG bundle is configured for one-to-one active/backup link redundancy. To support link redundancy at the MPC level, the LAG bundle attaches to ports from two different MPCs.

   ```
   [edit interfaces]
   user@host# set ge-5/0/3 gigether-options 802.3ad ae0
   user@host# set ge-5/0/3 gigether-options 802.3ad primary
   user@host# set ge-5/1/2 gigether-options 802.3ad ae0
   user@host# set ge-5/1/2 gigether-options 802.3ad backup
   ```

3. Enable link protection on the aggregated Ethernet logical interface and configure support for single and dual (stacked) VLAN tags.

   ```
   [edit interfaces]
   user@host# set ae0 aggregated-ether-options link-protection
   user@host# set ae0 flexible-vlan-tagging
   ```

4. Configure the parameters for automatically configuring VLANs and S-VLANs, including the VLAN ranges and dynamic profiles.

   ```
   [edit interfaces]
   user@host# set ae0 auto-configure vlan-ranges dynamic-profile vlan-profile accept
     pppoe
   ```


Copyright © 2013, Juniper Networks, Inc.

```
user@host# set ae0 auto-configure vlan-ranges dynamic-profile vlan-profile ranges
  1-4094
user@host# set ae0 auto-configure stacked-vlan-ranges dynamic-profile
  svlan-profile accept pppoe
user@host# set ae0 auto-configure stacked-vlan-ranges dynamic-profile
  svlan-profile ranges 1-4094,1-4094
```

5. Configure the dynamic profile that creates the PPPoE subscriber interface.

```
[edit dynamic-profiles pppoe-profile]
user@host# edit interfaces pp0 unit $junos-interface-unit
[edit dynamic-profiles pppoe-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options underlying-interface $junos-underlying-interface
user@host# set pppoe-options server
user@host# set family inet unnumbered-address lo0.0
```

6. Configure the dynamic profile that creates VLAN demux underlying interfaces, including the PPPoE family attributes.

```
[edit dynamic-profiles vlan-profile]
user@host# edit interfaces demux0 unit $junos-interface-unit
[edit dynamic-profiles vlan-profile interfaces demux0 unit "$junos-interface-unit"]
user@host# set vlan-id $junos-vlan-id
user@host# set demux-options underlying-interface $junos-interface-ifd-name
user@host# set family pppoe access-concentrator pppoe-server-1
user@host# set family pppoe duplicate-protection
user@host# set family pppoe dynamic-profile pppoe-profile
```

7. Configure the dynamic profile that creates S-VLAN demux underlying interfaces, including the PPPoE family attributes.

```
[edit dynamic-profiles svlan-profile]
user@host# edit interfaces demux0 unit $junos-interface-unit
[edit dynamic-profiles svlan-profile interfaces demux0 unit "$junos-interface-unit"]
user@host# set vlan-tags outer $junos-stacked-vlan-id
user@host# set vlan-tags inner $junos-vlan-id
user@host# set demux-options underlying-interface $junos-interface-ifd-name
user@host# set family pppoe access-concentrator pppoe-server-1
user@host# set family pppoe duplicate-protection
user@host# set family pppoe dynamic-profile pppoe-profile
```

Results From configuration mode, confirm the aggregated device configuration by entering the **show chassis** command. Confirm the interface configuration by entering the **show interfaces** command. Confirm the dynamic profile configuration by entering the **show dynamic-profiles** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show chassis
aggregated-devices {
   ethernet {
      device-count 1;
   }
}
```

```
[edit]
user@host# show interfaces
```

```
    ge-5/0/3 {
      gigether-options {
        802.3ad {
          ae0;
          primary;
        }
      }
    }
    ge-5/1/2 {
      gigether-options {
        802.3ad {
          ae0;
          backup;
        }
      }
    }
    ae0 {
      flexible-vlan-tagging;
      aggregated-ether-options {
        link-protection;
      }
      auto-configure {
        vlan-ranges {
          dynamic-profile {
            vlan-profile {
              accept pppoe;
              vlan-ranges 1–4094};
            }
          }
        }
        stacked-vlan-ranges {
          dynamic-profile {
            svlan-profile {
              accept pppoe;
              vlan-ranges 1–4094,1–4094;
            }
          }
        }
      }
    }

[edit]
user@host# show dynamic-profiles
pppoe-profile {
  interfaces {
    pp0 {
      unit $junos-interface-unit {
        pppoe-options {
          underlying-interface $junos-underlying-interface;
          server;
        }
        family inet {
          unnumbered-address lo0.0;
        }
      }
    }
```

 Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  vlan-profile {
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          vlan-id "$junos-vlan-id";
          demux-options {
            underlying-interface "$junos-interface-ifd-name";
          }
          family pppoe {
            access-concentrator pppoe-server-1;
            duplicate-protection;
            dynamic-profile pppoe-profile;
          }
        }
      }
    }
  }
  svlan-profile {
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          vlan-tags outer "$junos-stacked-vlan-id" inner "$junos-vlan-id";
          demux-options {
            underlying-interface "$junos-interface-ifd-name";
          }
          family pppoe {
            access-concentrator pppoe-server-1;
            duplicate-protection;
            dynamic-profile pppoe-profile;
          }
        }
      }
    }
  }
```

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

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying the Aggregated Ethernet Interface Configuration on page 881

Verifying the Aggregated Ethernet Interface Configuration

Purpose Verify that the interface values match your configuration, the link is up, and traffic is flowing.

Action From operational mode, enter the **show interfaces redundancy** command.

```
user@host> show interfaces redundancy
Interface  State           Last change  Primary      Secondary    Current status
ae0        On primary                   ge-5/0/3     ge-5/1/2     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: 00:1f:12:b8:ef:c0, Hardware address: 00:1f:12:b8:ef:c0
  Last flapped   : 2011-03-11 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.

Related Documentation

- Subscriber Interfaces and Demultiplexing Overview on page 759
- Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview on page 816
- Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles on page 772
- Configuring the PPPoE Family for an Underlying Interface on page 831
- Configuring a Basic PPPoE Dynamic Profile on page 900

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 46

Dynamic PPPoE Subscriber Interfaces Overview

- Subscriber Interfaces and PPPoE Overview on page 883
- Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview on page 887
- PPPoE Maximum Session Limit Overview on page 890
- Guidelines for Using PPPoE Maximum Session Limit from RADIUS on page 892
- PPPoE Subscriber Session Lockout Overview on page 893
- Understanding the Lockout Period for PPPoE Subscriber Session Lockout on page 896

Subscriber Interfaces and PPPoE Overview

You can configure the router to dynamically create Point-to-Point Protocol over Ethernet (PPPoE) logical interfaces on statically created underlying Ethernet interfaces. The router automatically and transparently creates the dynamic interface in response to the receipt of a PPPoE Active Discovery Request (PADR) control packet on the underlying interface. Because the router creates a dynamic PPPoE logical interface on demand when a subscriber logs in to the network, dynamic PPPoE logical interfaces are also referred to as *dynamic PPPoE subscriber interfaces*.

To enable the router to create a dynamic PPPoE logical interface on an underlying Ethernet interface, you define the attributes of the PPPoE logical interface in a dynamic profile, and then attach the dynamic profile to an Ethernet interface configured with PPPoE encapsulation. When the router receives a PADR control packet from a PPPoE client on an underlying interface with a PPPoE dynamic profile attached, the router uses the attributes defined in the profile to instantiate a dynamic PPPoE subscriber interface to handle the PPPoE session.

This overview covers the following topics:

- Benefits of Using Dynamic PPPoE Subscriber Interfaces on page 884
- Supported Platforms for Dynamic PPPoE Subscriber Interfaces on page 885
- Sequence of Operations for PPPoE Subscriber Access on page 885

Copyright © 2013, Juniper Networks, Inc.

Benefits of Using Dynamic PPPoE Subscriber Interfaces

Configuring and using dynamic PPPoE subscriber interfaces offers the following benefits:

- On-demand dynamic interface creation

  Configuring dynamic PPPoE subscriber interfaces provides the flexibility of dynamically creating the PPPoE subscriber interface only when needed; that is, when a subscriber logs in on the associated underlying Ethernet interface. By contrast, statically created interfaces allocate and consume system resources when the interface is created. Configuring and using dynamically created interfaces helps you effectively and conveniently manage edge or access networks in which large numbers of subscribers are constantly logging in to and logging out from the network on a transient basis.

- Dynamic removal of PPPoE subscriber interfaces without manual intervention

  When the PPPoE subscriber logs out or the PPPoE session is terminated, the router dynamically deletes the associated PPPoE subscriber interface without your intervention, thereby restoring any consumed resources to the router.

- Use of dynamic profiles to efficiently manage multiple subscriber interfaces

  A *dynamic profile* is a set of characteristics that can be dynamically assigned to subscriber interfaces. By using a profile, you reduce the management of a large number of interfaces by applying a set of common characteristics to multiple interfaces. When you configure a dynamic profile for PPPoE, you use predefined dynamic variables in the profile to represent information that varies from subscriber to subscriber, such as the logical unit number and underlying interface name. These variables are dynamically replaced with the values supplied by the network when the subscriber logs in.

- Denial of service (DoS) protection

  You can optionally configure the underlying Ethernet interface with certain PPPoE-specific attributes that can reduce the potential for DoS attacks. Duplicate protection, which is disabled by default, prevents activation of another dynamic PPPoE logical interface on the underlying interface when a PPPoE logical interface for the same client is already active on the underlying interface. You can also specify the maximum number of PPPoE sessions that the router can activate on the underlying interface. By enabling duplicate protection and restricting the maximum number of PPPoE sessions on the underlying interface, you can ensure that a single toxic PPPoE client cannot monopolize allocation of the PPPoE session.

- Support for dynamic PPPoE subscriber interface creation from PPPoE service name tables

  You can assign a previously configured PPPoE dynamic profile to a named, **empty**, or **any** service entry in a PPPoE service name table, or to an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services. The router uses the attributes defined in the profile to instantiate a dynamic PPPoE subscriber interface based on the service name, ACI, and ARI information provided by the PPPoE client during PPPoE negotiation. To specify the routing instance in which to instantiate the dynamic PPPoE subscriber interface, you can assign a previously configured routing instance to a named, **empty**, or **any** service, or to an ACI/ARI pair defined for these services. The dynamic

 Copyright © 2013, Juniper Networks, Inc.

profile and routing instance configured for the PPPoE service name table overrides the dynamic profile and routing instance assigned to the PPPoE underlying interface on which the dynamic subscriber interface is created.

Supported Platforms for Dynamic PPPoE Subscriber Interfaces

Configuration of dynamic PPPoE subscriber interfaces over static underlying Ethernet interfaces is supported on the following routing platforms:

- Intelligent Queuing 2 (IQ2) PICs on M120 Multiservice Edge Router and M320 Multiservice Edge Router
- MPC/MIC interfaces on MX Series 3D Universal Edge Routers

Sequence of Operations for PPPoE Subscriber Access

When a PPPoE subscriber logs in to the network, the PPPoE protocol defines the sequence of operations by which a connection is established and traffic flow is enabled on the dynamic PPPoE subscriber interface. Similarly, when the PPPoE subscriber logs out from the network, PPPoE defines the sequence that occurs to terminate the connection and remove the dynamic PPPoE subscriber interface from the router.

The router creates a dynamic PPPoE subscriber interface for each new PPPoE session, and removes the dynamic PPPoE subscriber interface when the session is terminated due to subscriber logout, PPP negotiation failure, or down status of the underlying Ethernet interface. Dynamic PPPoE subscriber interfaces are never reused for multiple PPPoE sessions.

Sequence When a PPPoE Subscriber Logs In

In a PPPoE subscriber network, the router acts as a *remote access concentrator*, also known as a *PPPoE server*. For a PPPoE client to initiate a PPPoE session with a PPPoE server, it must first perform PPPoE Discovery to identify the Ethernet MAC address of the remote access concentrator that can service its request. Based on the network topology, there may be more than one remote access concentrator with which the client can communicate. The Discovery process enables a PPPoE client to find all remote access concentrators and then select one to connect to.

The following sequence occurs when a PPPoE subscriber logs in to the network. Steps 1 through 5 in this sequence are part of the PPPoE Discovery process.

1. The PPPoE client broadcasts a PPPoE Active Discovery Initiation (PADI) packet to all remote access concentrators in the network.
2. One or more remote access concentrators respond to the PADI packet by sending a PPPoE Active Discovery Offer (PADO) packet, indicating that they can service the client request. The PADO packet includes the name of the access concentrator from which it was sent.
3. The client sends a unicast PPPoE Active Discovery Request (PADR) packet to the access concentrator it selects.

Copyright © 2013, Juniper Networks, Inc.

4. On receipt of the PADR packet on the underlying interface associated with a PPPoE dynamic profile, the router uses the attributes configured in the dynamic profile to create the dynamic PPPoE logical interface.
5. The router sends a PPPoE Active Discovery Session (PADS) packet to confirm establishment of the PPPoE connection.
6. The PPP Link Control Protocol (LCP) negotiates the PPP link between the client and the PPPoE server.
7. The subscriber is authenticated using the PPP authentication protocol (CHAP or PAP) configured in the PPPoE dynamic profile.
8. The PPP Network Control Protocol (NCP) negotiates the IP routing protocol and network family.
9. The PPP server issues an IP access address for the client, and the router adds the client access route to its routing table.
10. The router instantiates the dynamic profile and applies the attributes configured in the profile to the dynamic PPPoE subscriber interface.
11. PPP NCP negotiation completes, enabling traffic flow between the PPPoE client and the PPPoE server.

Sequence When a PPPoE Subscriber Logs Out

The following sequence occurs when a PPPoE subscriber logs out of the network:

1. The client terminates the PPP connection and the router receives an LCP termination request.
2. The router removes the client access router from its routing table.
3. The router sends or receives a PPPoE Active Discovery Termination (PADT) packet to end the PPPoE connection.
4. The router deactivates the subscriber, gathers final statistics for the PPPoE session, and sends the RADIUS server an Acct-Stop accounting message.
5. The router de-instantiates the PPPoE dynamic profile and removes the PPPoE logical interface. The router does not reuse the PPPoE logical interface for future dynamic PPPoE sessions.

Related Documentation

- Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview on page 887
- Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles on page 899
- For information about configuring static PPPoE interfaces and PPPoE service name tables, see the *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

Creating a dynamic PPPoE subscriber interface over a static underlying Ethernet interface consists of two basic steps:

1. Configure a dynamic profile to define the attributes of the PPPoE logical interface.
2. Attach the dynamic profile to a statically created underlying Ethernet interface configured with PPPoE encapsulation.

This overview describes the concepts you need to understand to configure a dynamic PPPoE subscriber interface, and covers the following topics:

- PPPoE Dynamic Profile Configuration on page 887
- PPPoE Underlying Interface Configuration on page 888
- Address Assignment for Dynamic PPPoE Subscriber Interfaces on page 888
- Guidelines for Configuring Dynamic PPPoE Subscriber Interfaces on page 889

PPPoE Dynamic Profile Configuration

A *dynamic profile* is a template for configuring a dynamic interface. You use predefined dynamic variables in the PPPoE dynamic profile to represent information that varies from subscriber to subscriber, such as the logical unit number and underlying interface name. These variables are dynamically replaced with the values supplied by the network when the subscriber logs in. On receipt of traffic on an underlying Ethernet interface to which a dynamic profile is attached, the router creates the dynamic PPPoE logical interface, also referred to as a *dynamic PPPoE subscriber interface*, on the underlying interface and applies the properties configured in the dynamic profile.

To provide basic access for PPPoE subscribers, the dynamic profile must provide a minimal configuration for a **pp0** (PPPoE) logical interface that includes at least the following attributes:

- The logical unit number, represented by the **$junos-interface-unit** predefined dynamic variable
- The name of the underlying Ethernet interface, represented by the **$junos-underlying-interface** predefined dynamic variable
- Configuration of the router to act as a PPPoE server
- The PPP authentication protocol (PAP or CHAP)
- The unnumbered address for the **inet** (IPv4) or **inet6** (IPv6) protocol family

You can also optionally configure additional options for PPPoE subscriber access in the dynamic profile, including:

- The keepalive interval, or the option to disable sending keepalive messages
- The IPv4 or IPv6 address of the dynamic PPPoE logical interface

Copyright © 2013, Juniper Networks, Inc.

- The service sets and filters, input filters, and output filters to be applied to the dynamic PPPoE logical interface

PPPoE Underlying Interface Configuration

After you configure a dynamic profile to define the attributes of a dynamic PPPoE subscriber interface, you must attach the dynamic profile to the underlying Ethernet interface on which you want the router to dynamically create the PPPoE logical interface. The underlying interface for a dynamic PPPoE logical interface must be statically created and configured with PPPoE (**ppp-over-ether**) encapsulation. When a PPPoE subscriber logs in on the underlying interface, the router dynamically creates the PPPoE logical interface and applies the attributes defined in the profile to the interface.

In addition to attaching the dynamic profile to the interface, you can also configure the underlying interface with one or more of the following optional PPPoE-specific attributes:

- Prevention of another dynamic PPPoE logical interface from being activated on the underlying interface when a PPPoE logical interface for a client with the same MAC address is already active on that interface
- Maximum number of dynamic PPPoE logical interfaces (sessions) that the router can activate on the underlying interface
- An alternative access concentrator name in the AC-NAME tag in a PPPoE control packet

Address Assignment for Dynamic PPPoE Subscriber Interfaces

If the subscriber address for a dynamic PPPoE interface is not specified by means of the Framed-IP-Address (8) or Framed-Pool (88) RADIUS IETF attributes during authentication, the router allocates an IP address from the first IPv4 local address-assignment pool defined in the routing instance. For this reason, make sure that the local address assigned for the **inet** (IPv4) address family is in the same subnet as the addresses obtained from the first IPv4 local address-assignment pool.

The router allocates the IP address from the first IPv4 local address-assignment pool under either of the following conditions:

- RADIUS returns no address attributes.
- RADIUS authentication does not take place because only address allocation is requested.

If the first IPv4 local address-assignment pool has no available addresses, or if no IPv4 local address-assignment pools are configured, the router does not allocate an IP address to the dynamic PPPoE subscriber interface, and denies access to the associated subscriber. To avoid depletion of IP addresses, you can configure linked address-assignment pools on the first IPv4 local address-assignment pool to create one or more backup pools.

For more information, see "Configuring Address-Assignment Pools" on page 168.

 Copyright © 2013, Juniper Networks, Inc.

Guidelines for Configuring Dynamic PPPoE Subscriber Interfaces

Observe the following guidelines when you configure dynamic PPPoE subscriber interfaces:

- You can configure dynamic PPPoE subscriber interfaces for the **inet** (IPv4) and **inet6** (IPv6) protocol families.
- When you configure the **pp0** (PPPoE) logical interface in a PPPoE dynamic profile, you must include the **pppoe-options** subhierarchy at the **[edit dynamic-profiles** ***profile-name*** **interfaces pp0 unit "$junos-interface-unit"]** hierarchy level. At a minimum, the **pppoe-options** subhierarchy must include the name of the underlying Ethernet interface, represented by the **$junos-underlying-interface** predefined dynamic variable, and the **server** statement, which configures the router to act as a PPPoE server. If you omit the **pppoe-options** subhierarchy from the configuration, the **commit** operation fails.
- When you configure CHAP or PAP authentication in a PPPoE dynamic profile, you cannot configure additional options for the **chap** or **pap** statements. This is because the router supports only unidirectional authentication for dynamic interfaces; that is, the router always functions as the authenticator.
- When you attach the PPPoE dynamic profile to an underlying Ethernet interface, ensure that both of the following conditions are met:
  - The PPPoE dynamic profile has already been configured on the router.
  - The underlying Ethernet interface has already been statically configured on the router with PPPoE (**ppp-over-ether**) encapsulation.
- You cannot attach a PPPoE dynamic profile to an underlying Ethernet interface that is already associated with static PPPoE logical interfaces. Conversely, you cannot associate static PPPoE logical interfaces with an underlying Ethernet interface that already has a PPPoE dynamic profile attached.

Related Documentation

- Subscriber Interfaces and PPPoE Overview on page 883
- Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles on page 899
- Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Gigabit Ethernet VLAN Interface on page 917
- For more information about static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

PPPoE Maximum Session Limit Overview

The maximum session limit for PPPoE subscriber interfaces specifies the maximum number of concurrent static or dynamic PPPoE logical interfaces (sessions) that the router can activate on the PPPoE underlying interface, or the maximum number of active static or dynamic PPPoE sessions that the router can establish with a particular service entry in a PPPoE service name table.

You can configure the PPPoE maximum session limit in one of two ways:

- On a per-interface basis, by using the **max-sessions** CLI statement
- (Default) On a per-subscriber basis, by using the value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks vendor-specific attribute (VSA) [26-143]

This overview describes the concepts you need to understand to configure the PPPoE maximum session limit, and covers the following topics:

- Per-Interface Configuration for PPPoE Maximum Session Limit Using the CLI on page 890
- Per-Subscriber Configuration for PPPoE Maximum Session Limit Using RADIUS on page 890
- Override of PPPoE Maximum Session Limit from RADIUS on page 891

Per-Interface Configuration for PPPoE Maximum Session Limit Using the CLI

To configure the PPPoE maximum session limit for a particular interface, you can use the **max-sessions** statement to specify either or both of the following, depending on the hierarchy level at which you include the statement:

- The maximum number of concurrent PPPoE sessions that the router can activate on the PPPoE underlying interface
- The maximum number of active PPPoE sessions using either static or dynamic PPPoE interfaces that the router can establish with a particular named service entry, **empty** service entry, or **any** service entry in a PPPoE service name table

You can configure the PPPoE maximum session value from 1 through the platform-specific default for your router. The default value is equal to the maximum number of PPPoE sessions supported on your routing platform. If the number of active PPPoE sessions exceeds the value configured with the **max-sessions** statement, the router prohibits creation of any new PPPoE sessions, and the PPPoE application on the router returns a PPPoE Active Discovery Session (PADS) packet with an error to the PPPoE client.

Changing the PPPoE maximum session value has no effect on dynamic PPPoE subscriber interfaces that are already active.

Per-Subscriber Configuration for PPPoE Maximum Session Limit Using RADIUS

To configure the PPPoE maximum session limit for a particular subscriber, you can use the value returned by the RADIUS server in the Max-Clients-Per-Interface Juniper Networks VSA [26-143] during the subscriber authentication process. For PPPoE clients, the

 Copyright © 2013, Juniper Networks, Inc.

Max-Clients-Per-Interface VSA returns the maximum number of sessions (PPPoE subinterfaces) per PPPoE major interface.

By default, the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA takes precedence over the PPPoE maximum session value configured with the **max-sessions** statement.

If you configure multiple subscribers on the same PPPoE underlying VLAN interface and RADIUS returns a different PPPoE maximum session value for each subscriber, the router uses the most recent PPPoE maximum session value returned by RADIUS to determine whether to override the current PPPoE maximum session value and create the new PPPoE session.

The following sequence describes how the router obtains the PPPoE maximum session value from RADIUS when a PPPoE subscriber logs in to initiate a session with the router. (In a PPPoE subscriber network, the router functions as a *remote access concentrator*, also known as a *PPPoE server*.)

1. The PPPoE client and the router participate in the PPPoE Discovery process to establish the PPPoE connection.
2. The PPP Link Control Protocol (LCP) negotiates the PPP link between the client and the router.
3. The PPP application sends the subscriber authentication request to the AAA application.
4. AAA sends the authentication request to an external RADIUS server.
5. The RADIUS server returns the PPPoE maximum session value for that subscriber to AAA in the Max-Clients-Per-Interface VSA as part of an Access-Accept message.

   NOTE: **The RADIUS server does not return the Max-Clients-Per-Interface VSA in Change of Authorization Request (CoA-Request) messages.**

6. AAA passes the response from RADIUS to PPP.
7. PPP validates the subscriber parameters and, if authentication succeeds, passes the PPPoE maximum session value returned by RADIUS to the PPPoE application.
8. PPPoE uses the maximum session value returned by RADIUS to determine whether to override the current PPPoE maximum session value and create or tear down the new PPPoE session.

Override of PPPoE Maximum Session Limit from RADIUS

By default, the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA [26-143] takes precedence over the PPPoE maximum session value configured with the **max-sessions** statement. To configure the router to ignore (clear) the PPPoE maximum session value returned by the RADIUS server in the Max-Clients-Per-Interface VSA, include the **max-sessions-vsa-ignore** statement at the same hierarchy levels that you can specify the **max-sessions** statement.

Including the **max-sessions-vsa-ignore** statement in your configuration restores the PPPoE maximum session value on the underlying interface to the value configured in the CLI with the **max-sessions** statement.

Related Documentation

- Guidelines for Using PPPoE Maximum Session Limit from RADIUS on page 892
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface on page 908
- For more information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

Guidelines for Using PPPoE Maximum Session Limit from RADIUS

Consider the following guidelines when you use the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface vendor-specific attribute (VSA) [26-143]:

- If the current number of sessions (including newly created sessions) is *less than* the new PPPoE maximum session value returned by RADIUS, the PPPoE application overrides the current value and enables interface creation to proceed.
- If the current number of sessions (including newly created sessions) is *equal to* the new PPPoE maximum session value returned by RADIUS, the PPPoE application overrides the current value and enables interface creation to proceed.
- If the current number of sessions (including newly created sessions) is *greater than* the new PPPoE maximum session value returned by RADIUS, the PPPoE application overrides the current value and brings down the new interface.

To illustrate these guidelines, Table 70 on page 892 shows examples of how the router handles the PPPoE session when the current number of sessions is less than (first row), equal to (second row), and greater than (third row) the new PPPoE maximum session value returned by RADIUS when a new subscriber logs in.

Table 70: Sample PPPoE Maximum Session Values During Subscriber Login

New PPPoE Maximum Session Value from RADIUS	Current PPPoE Maximum Session Value	Existing Number of PPPoE Sessions	New PPPoE Maximum Session Value	New Number of PPPoE Sessions	Status of Session
10	5	4	10	5	PPPoE session up
5	5	4	5	5	PPPoE session up
3	5	4	3	4	PPPoE session down

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- PPPoE Maximum Session Limit Overview on page 890
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface on page 908
- For more information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

PPPoE Subscriber Session Lockout Overview

PPPoE subscriber session lockout, which is sometimes referred to as PPPoE encapsulation type lockout, configures the router to temporarily prevent (lock out) a failed or short-lived static or dynamic PPPoE subscriber session from reconnecting for a default or configurable period of time. This time period, known as the *lockout period*, is derived from a formula and increases exponentially based on the number of successive reconnection failures.

This overview describes the concepts you need to understand to configure PPPoE subscriber session lockout, and covers the following topics:

- Benefits of Using PPPoE Subscriber Session Lockout on page 893
- Supported Platforms and Underlying Interfaces for PPPoE Subscriber Session Lockout on page 894
- How PPPoE Subscriber Session Lockout Works on page 894
- PPPoE Subscriber Session Lockout Period on page 895
- PPPoE Subscriber Session Lockout and Duplicate Protection on page 895
- PPPoE Subscriber Session Lockout and Automatic Removal of Dynamic Subscriber VLANs on page 896

Benefits of Using PPPoE Subscriber Session Lockout

Configuring and using PPPoE subscriber session lockout provides the following benefits:

- Reduces excessive loading on the router

  By temporarily locking out failed or short-lived PPPoE sessions, PPPoE subscriber session lockout protects the router from excessive loading by:

  - Reducing the resources required to receive and process PPPoE control packets to negotiate and terminate short-lived connections
  - Reducing the resources required to allocate and deallocate services, such as class of service (CoS) and firewall filters, for failed or short-lived subscriber sessions

  PPPoE subscriber session lockout increases router efficiency by temporarily deferring failed or short-lived subscriber sessions in favor of those sessions that can complete successfully.

- Reduces excessive loading on external authentication, authorization, and accounting (AAA) servers

PPPoE subscriber session lockout protects any external AAA servers, such as RADIUS or Diameter, from excessive loading:

  - As a result of failed or short-lived PPPoE subscriber sessions that occur repeatedly for the same subscriber
  - By reducing the resources required to authenticate and terminate these connections

- Enables lockout of a single failed or short-lived PPP session without disrupting other PPP sessions on the same PPPoE underlying interface

  In some subscriber network configurations, the PPPoE underlying interface supports multiple upper-layer PPP sessions. Because PPPoE subscriber session lockout identifies each subscriber session by its unique media access control (MAC) source address on the underlying interface, the router is able to lock out only the offending PPP session while enabling other PPP sessions on the same underlying interface to successfully negotiate the connection.

Supported Platforms and Underlying Interfaces for PPPoE Subscriber Session Lockout

You can configure PPPoE subscriber session lockout on the following platforms and underlying interface types:

- Supported platforms:
  - Intelligent Queuing 2 (IQ2) PICs on M120 Multiservice Edge Router and M320 Multiservice Edge Router
  - MPC/MIC interfaces on MX Series 3D Universal Edge Routers
- Supported PPPoE underlying interfaces:
  - Static VLAN logical interface
  - Static VLAN demultiplexing (demux) logical interface
  - Dynamic VLAN logical interface
  - Dynamic VLAN demux logical interface

How PPPoE Subscriber Session Lockout Works

PPPoE subscriber session lockout is disabled on the router by default. When you enable PPPoE subscriber session lockout by issuing the **short-cycle-protection** statement, the router does the following:

1. Detects a short-lived subscriber session, also referred to as a *short-cycle event*.

   A short-lived subscriber session is detected, partially or completely created, and terminated by the router within 150 seconds. The router identifies each PPPoE subscriber session by its unique MAC source address on the PPPoE underlying interface.

2. Tracks the time between repeated short-cycle events to determine whether to increase the lockout time for a subsequent short-cycle event.

 Copyright © 2013, Juniper Networks, Inc.

3. Applies a time penalty for each short-cycle event based on a default or configured lockout period and the number of consecutive short-cycle events that occur repeatedly for the same subscriber.

   If you enable PPPoE subscriber session lockout but do not configure a lockout time range, the router uses the default lockout time range of 1 through 300 seconds (5 minutes).

4. Temporarily locks out the specified PPPoE subscriber by preventing connection to the router.

   During lockout, the router drops negotiation packets for the PPPoE subscriber session until the lockout period expires. When the lockout period expires, the PPPoE subscriber session and its associated MAC source address resume normal negotiation of the connection.

Repeated creation of multiple short-lived (short-cycle) PPPoE subscriber sessions can cause excessive loading on the router. Conditions that can cause a short-lived subscriber session include:

- Authentication denials from external AAA servers, such as RADIUS, due to the absence of a corresponding entry in the RADIUS database or due to improper login attempts
- Configuration errors within a dynamic profile or RADIUS record
- Insufficient memory resources to create a dynamic PPPoE subscriber interface
- Protocol failure or error within the dynamic PPPoE subscriber interface
- Client logout shortly after a successful login; this action creates a complete dynamic PPPoE subscriber interface before the interface is torn down

PPPoE Subscriber Session Lockout Period

The lockout period is the time during which the router temporarily prevents (locks out) a failed or short-lived PPPoE subscriber session identified by a unique MAC source address from reconnecting to the router. You can use the default lockout time range of 1 through 300 seconds (5 minutes), or you can override the default lockout period by configuring a nondefault lockout time in the range 1 through 86,400 seconds (24 hours).

PPPoE Subscriber Session Lockout and Duplicate Protection

Duplicate protection, which is disabled on the router by default, prevents the activation of another PPPoE subscriber session on the same PPPoE underlying interface when a PPPoE subscriber session with the same media access control (MAC) address is already active on that interface. When you configure PPPoE subscriber session lockout, we recommend that you enable duplicate protection to ensure that the MAC source address for each active PPPoE session is unique on the underlying interface.

With PPPoE subscriber session lockout configured, the router identifies subscriber sessions by their unique MAC source address. If the router detects a short-lived (short-cycle) subscriber session, it applies the default or configured lockout period to that MAC source address to temporarily prevent reconnection. If the MAC source address is not unique on

the underlying interface, multiple PPPoE subscriber sessions with the same MAC source address might also be affected by the lockout.

PPPoE Subscriber Session Lockout and Automatic Removal of Dynamic Subscriber VLANs

You can configure automatic removal of subscriber VLANs that have no PPPoE client sessions by issuing the **remove-when-no-subscribers** statement at the **[edit interfaces *interface-name* auto-configure]** hierarchy level. If PPPoE subscriber session lockout is also configured, the router does not remove the unused subscriber VLAN until the lockout time has expired for each PPPoE client undergoing lockout on the underlying interface.

Related Documentation

- Understanding the Lockout Period for PPPoE Subscriber Session Lockout on page 896
- Configuring Lockout of PPPoE Subscriber Sessions on page 912
- Clearing Lockout of PPPoE Subscriber Sessions on page 913
- Verifying and Managing Dynamic PPPoE Configuration on page 914
- For more information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

Understanding the Lockout Period for PPPoE Subscriber Session Lockout

When you configure PPPoE subscriber session lockout, the router applies a time penalty called the *lockout period* for each failed or short-lived subscriber session. During the lockout period, the router temporarily prevents (locks out) a failed or short-lived PPPoE subscriber session identified by a unique media access control (MAC) source address from reconnecting to the router.

This overview describes how the router determines and applies the PPPoE subscriber session lockout period, and covers the following topics:

- Duration of PPPoE Subscriber Session Lockout Period on page 896
- How the Router Determines the PPPoE Subscriber Session Lockout Period on page 897

Duration of PPPoE Subscriber Session Lockout Period

The duration of the lockout period is based on a default or configured lockout time and the number of consecutive short-cycle (short-lived) events that occur repeatedly for the same subscriber. When you include the **short-cycle-protection** statement to configure PPPoE subscriber session lockout on a PPPoE underlying interface, you can use the default lockout time range of 1 through 300 seconds (5 minutes), or you can override the default lockout period by configuring a nondefault lockout time in the range 1 through 86,400 seconds (24 hours).

The lockout time penalty applied by the router for each short-cycle event differs depending on the event. For example, some short-cycle events represent normal subscriber behavior, such as a PPPoE subscriber logging in once per hour to check e-mail and logging out shortly thereafter. The router does not noticeably penalize a subscriber for these types of events.

 Copyright © 2013, Juniper Networks, Inc.

By contrast, other short-cycle events are the result of repeated attempts to log in to the router for reasons such as an incorrectly typed password, customer premises equipment (CPE) that performs repeated auto-retries, or malicious attempts to access the Internet illegally. For these types of short-cycle events, the router applies a lockout time penalty that starts with a short time interval and increases exponentially. In these instances, the initial lockout time is short enough to avoid noticeably penalizing a subscriber who, for example, types a password incorrectly several times before entering the correct one.

For example, using the default lockout time range of 1 through 300 seconds, the increasing lockout period on the router is: 1 second, 2 seconds, 4 seconds, 8 seconds, 16 seconds, 32 seconds, 64 seconds, 128 seconds, 256 seconds, and finally, 300 seconds (5 minutes).

How the Router Determines the PPPoE Subscriber Session Lockout Period

The router uses the following rules to determine the PPPoE subscriber session lockout period for short-lived PPPoE subscriber sessions:

- The lockout period is derived from the following formula:

  (minimum lockout time) * (2 ^ *n*-1)

  where *n* represents the number of consecutive short-cycle events for the same subscriber. The router identifies a PPPoE subscriber session by its MAC source address, which should be unique on the underlying PPPoE interface.

- The router increments the value of *n* when the time between short-cycle events is either within 15 minutes or the maximum lockout time, whichever is greater.

- When the time between short-cycle events is greater than either 15 minutes or the maximum lockout time, the value of *n* reverts to 1. This condition is referred to as a *lockout grace period*.

- The lockout period never exceeds the maximum configured lockout time.

  For example, for a configured (nondefault) lockout time in the range 20 through 120 seconds, the increasing lockout period on the router is: 20 seconds, 40 seconds, 80 seconds, and finally, 120 seconds (2 minutes).

- A *short-cycle event* is detected, partially or completely created, and terminated by the router within 150 seconds. The router tracks the time between short-cycle events to determine whether to increase the lockout time for a subsequent short-cycle event for the same subscriber.

  NOTE: **When the calculated lockout time is equal to or exceeds the maximum lockout time, the router uses the maximum lockout time value until the time to the next short-cycle event exceeds the greater of 15 minutes or the maximum lockout time value. At that point, the lockout time reverts to the minimum lockout time value.**

- The minimum lockout time value cannot exceed the maximum lockout time value.

  When the minimum and maximum lockout time values are equal, the lockout time becomes fixed at that value.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- PPPoE Subscriber Session Lockout Overview on page 893
- Configuring Lockout of PPPoE Subscriber Sessions on page 912
- Clearing Lockout of PPPoE Subscriber Sessions on page 913
- Verifying and Managing Dynamic PPPoE Configuration on page 914
- For more information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 47

Configuring Dynamic PPPoE Subscriber Interfaces

- Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles on page 899
- Configuring a Basic PPPoE Dynamic Profile on page 900
- Configuring a PPPoE Dynamic Profile with Additional Options on page 903
- Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces on page 905
- Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface on page 908
- Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation on page 910
- Configuring Lockout of PPPoE Subscriber Sessions on page 912
- Clearing Lockout of PPPoE Subscriber Sessions on page 913
- Verifying and Managing Dynamic PPPoE Configuration on page 914

Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles

You can configure dynamic PPP-over-Ethernet (PPPoE) subscriber interfaces by using dynamic profiles. To enable the router to create a dynamic PPPoE subscriber interface on a PPPoE underlying interface, you define the attributes of the PPPoE logical interface in a dynamic profile, and then configure the underlying interface to use the dynamic profile.

To configure a dynamic PPPoE subscriber interface:

1. Configure a dynamic profile to define the attributes of the PPPoE logical interface.
   - To configure a basic dynamic profile for PPPoE subscriber access, see "Configuring a Basic PPPoE Dynamic Profile" on page 900.
   - To configure a dynamic profile for PPPoE with additional options for subscriber access, see "Configuring a PPPoE Dynamic Profile with Additional Options" on page 903.
2. Configure the underlying Ethernet interface to use the dynamic profile for PPPoE.

   See "Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces" on page 905.

Copyright © 2013, Juniper Networks, Inc.

3. (Optional) Assign a dynamic profile and routing instance to a service name or ACI/ARI pair in a PPPoE service name table to instantiate a dynamic PPPoE subscriber interface based on the information provided by the PPPoE client.

   See "Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation" on page 910.

4. (Optional) Verify the dynamic PPPoE configuration by displaying or clearing PPPoE session statistics, and displaying information about the underlying Ethernet interface and PPPoE logical interface.

   See "Verifying and Managing Dynamic PPPoE Configuration" on page 914.

Related Documentation

- Subscriber Interfaces and PPPoE Overview on page 883
- Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview on page 887
- Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Gigabit Ethernet VLAN Interface on page 917
- Example: Configuring a PPPoE Service Name Table for Dynamic Subscriber Interface Creation on page 919

Configuring a Basic PPPoE Dynamic Profile

You can configure a basic dynamic profile for PPPoE subscribers that access the network. The dynamic profile defines the attributes of the dynamic PPPoE logical interface, also referred to as a *dynamic PPPoE subscriber interface*.

To provide basic access for PPPoE subscribers, the dynamic profile must provide a minimal configuration for a **pp0** (PPPoE) logical interface that includes the following:

- The logical unit number, represented by the **$junos-interface-unit** predefined dynamic variable
- The name of the underlying Ethernet interface, represented by the **$junos-underlying-interface** predefined dynamic variable
- The **server** statement, which configures the router to act as a PPPoE server
- The PPP authentication protocol (PAP or CHAP)
- The unnumbered address for the **inet** (IPv4) or **inet6** (IPv6) protocol family

NOTE: **The creation of dynamic PPPoE subscriber interfaces is supported for the inet and inet6 protocol families.**

To configure a basic PPPoE dynamic profile:

1. Name the dynamic profile.

   [edit]

 Copyright © 2013, Juniper Networks, Inc.

```
user@host# edit dynamic-profiles basic-pppoe-profile
```

2. Specify that you want to configure the **pp0** logical interface in the dynamic profile.

```
[edit dynamic-profiles basic-pppoe-profile]
user@host# edit interfaces pp0
```

3. Configure the predefined variable to represent the logical unit number for the **pp0** interface.

   You must use the **$junos-interface-unit** variable instead of the logical unit number for the **unit** statement. The **$junos-interface-unit** variable is dynamically replaced with the actual unit number supplied by the network when the subscriber logs in.

```
[edit dynamic-profiles basic-pppoe-profile interfaces pp0]
user@host# edit unit $junos-interface-unit
```

4. Configure PPPoE-specific options for the **pp0** interface.

   a. Configure the predefined variable to represent the name of the underlying Ethernet interface on which the router creates the dynamic PPPoE logical interface.

      You must use the **$junos-underlying-interface** variable instead of the underlying interface name for the **underlying-interface** statement. The **$junos-underlying-interface** variable is dynamically replaced with the actual name of the underlying interface supplied by the network when the subscriber logs in.

```
[edit dynamic-profiles basic-pppoe-profile interfaces pp0 unit
   "$junos-interface-unit"]
user@host# set pppoe-options underlying-interface $junos-underlying-interface
```

   b. Configure the router to act as a PPPoE server, also known as a remote access concentrator.

```
[edit dynamic-profiles basic-pppoe-profile interfaces pp0 unit
   "$junos-interface-unit"]
user@host# set pppoe-options server
```

5. Configure the PPP authentication protocol for the **pp0** interface.

   For dynamic interfaces, the router supports only unidirectional authentication; that is, the router always functions as the authenticator. When you configure PPP authentication in a dynamic profile, the **chap** and **pap** statements do not support any additional configuration options.

   - To configure CHAP authentication:

```
[edit dynamic-profiles basic-pppoe-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set ppp-options chap
```

   - To configure PAP authentication:

```
[edit dynamic-profiles basic-pppoe-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set ppp-options pap
```

6. Configure the protocol family for the **pp0** interface.

   a. Specify that you want to configure the **inet** (IPv4) or **inet6** (IPv6) protocol family.

```
[edit dynamic-profiles basic-pppoe-profile interfaces pp0 unit
   "$junos-interface-unit"]
```

```
user@host# edit family inet
```

b. Configure the unnumbered address for the protocol family.

```
[edit dynamic-profiles basic-pppoe-profile interfaces pp0 unit "$junos-interface-unit"
  family inet]
user@host# set unnumbered-address lo0.0
```

Related Documentation

- Subscriber Interfaces and PPPoE Overview on page 883
- Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview on page 887
- Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces on page 905
- Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Gigabit Ethernet VLAN Interface on page 917
- Verifying and Managing Dynamic PPPoE Configuration on page 914

Copyright © 2013, Juniper Networks, Inc.

Configuring a PPPoE Dynamic Profile with Additional Options

You can configure a dynamic profile for PPPoE that has additional options for subscriber access. All of these optional statements, with the exception of the **keepalives** and **no-keepalives** statements, are configured at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family*]** hierarchy level. The **keepalives** and **no-keepalives** statements are configured at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit"]** hierarchy level.

The additional options for PPPoE subscriber access in a dynamic profile can include one or more of the following:

- The keepalive interval (**keepalives**), or the option to disable sending keepalive messages (**no-keepalives**)
- The IPv4 or IPv6 address of the dynamic PPPoE logical interface (**address**)
- Definition of the service sets and filters to be applied to the dynamic PPPoE logical interface, configured at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* service]** hierarchy level
- Association of an input and output filter to the dynamic PPPoE logical interface, configured at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* filter]** hierarchy level

NOTE: **The creation of dynamic PPPoE subscriber interfaces is supported for the inet and inet6 protocol families.**

Before you begin:

- Configure a basic PPPoE dynamic profile.

  See "Configuring a Basic PPPoE Dynamic Profile" on page 900.

To configure a PPPoE dynamic profile with additional options for subscriber access:

1. Modify the keepalive interval, or configure the router to disable sending keepalive messages.
   - To modify the keepalive interval:

     ```
     [edit dynamic-profiles business-pppoe-profile interfaces pp0 unit
       "$junos-interface-unit"]
     user@host# set keepalives interval 15
     ```

   - To disable sending keepalive messages:

     ```
     [edit dynamic-profiles business-pppoe-profile interfaces pp0 unit
       "$junos-interface-unit"]
     user@host# set no-keepalives
     ```

2. Specify that you want to configure the **inet** (IPv4) or **inet6** (IPv6) protocol family.

```
[edit dynamic-profiles business-pppoe-profile interfaces pp0 unit
   "$junos-interface-unit"]
user@host# edit family inet
```

3. Specify the IPv4 or IPv6 address of the dynamic PPPoE logical interface.

```
[edit dynamic-profiles business-pppoe-profile interfaces pp0 unit
   "$junos-interface-unit" family inet]
user@host# set address 6.6.6.7/32
```

4. Specify the input and output service sets that you want to apply to the dynamic PPPoE logical interface.

```
[edit dynamic-profiles business-pppoe-profile interfaces pp0 unit
   "$junos-interface-unit" family inet]
user@host# set service input service-set inputService_100
user@host# set service input post-service-filter postService_20
user@host# set service output service-set outputService_200
```

5. Specify the input and output filters that you want to apply to the dynamic PPPoE logical interface.

   To control the order in which filters are processed, you can optionally specify a precedence value for the input filter, output filter, or both.

```
[edit dynamic-profiles business-pppoe-profile interfaces pp0 unit
   "$junos-interface-unit" family inet]
user@host# set filter input pppoe-input-filter
user@host# set filter output pppoe-output-filter precedence 50
```

Related Documentation

- Subscriber Interfaces and PPPoE Overview on page 883
- Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview on page 887
- Configuring a Basic PPPoE Dynamic Profile on page 900
- Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces on page 905
- Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Gigabit Ethernet VLAN Interface on page 917
- Verifying and Managing Dynamic PPPoE Configuration on page 914
- Dynamic Service Sets Overview on page 1151
- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197
- Dynamic Firewall Filters Overview on page 1132
- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type on page 1169

Copyright © 2013, Juniper Networks, Inc.

Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces

After you configure a dynamic profile to define the attributes of a dynamic PPPoE subscriber interface, you must attach the dynamic profile to a statically created underlying Ethernet interface configured with PPPoE (**ppp-over-ether**) encapsulation. You configure the underlying interface at the **[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]** hierarchy level.

In addition to attaching the dynamic profile to the interface by using the required **dynamic-profile** statement, you can also configure the underlying interface with one or more of the following optional PPPoE-specific attributes:

- Prevention of another dynamic PPPoE logical interface from being activated on the underlying interface when a PPPoE logical interface for a client with the same MAC address is already active on that interface (**duplicate-protection**)
- Maximum number of dynamic PPPoE logical interfaces (sessions) that the router can activate on the underlying interface on a per-interface basis (using the **max-sessions** statement) or on a per-subscriber basis (using the Max-Clients-Per-Interface Juniper Networks VSA [26-143])
- Lockout of failed or short-lived (also known as short-cycle) PPPoE subscriber sessions to prevent reconnection to the router for a default or configurable period of time (**short-cycle-protection**)
- An alternative access concentrator name in the AC-NAME tag in a PPPoE control packet (**access-concentrator**)

Before you begin:

1. Configure the static underlying Ethernet interface on which you want the router to dynamically create the PPPoE logical interface.

   For information about configuring static Ethernet interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

2. Configure a PPPoE dynamic profile in either of the following ways:

   - To configure a basic PPPoE dynamic profile, see "Configuring a Basic PPPoE Dynamic Profile" on page 900.
   - To configure a PPPoE dynamic profile with additional options for subscriber access, see "Configuring a PPPoE Dynamic Profile with Additional Options" on page 903.

To configure an underlying Ethernet interface for a dynamic PPPoE subscriber interface:

1. Specify the name and logical unit number of the static underlying Ethernet interface to which you want to attach the PPPoE dynamic profile.

   ```
   [edit]
   user@host# edit interfaces ge-1/0/1 unit 0
   ```

2. Configure PPPoE encapsulation on the underlying interface.

   ```
   [edit interfaces ge-1/0/1 unit 0]
   ```

```
user@host# set encapsulation ppp-over-ether
```

3. Specify that you want to configure PPPoE-specific options on the underlying interface.

```
[edit interfaces ge-1/0/1 unit 0]
user@host# edit pppoe-underlying-options
```

4. Attach a previously configured PPPoE dynamic profile to the underlying interface.

   The specified PPPoE dynamic profile must already be configured on the router. In addition, you cannot attach a PPPoE dynamic profile to an underlying Ethernet interface that is already associated with static PPPoE logical interfaces. Conversely, you cannot associate static PPPoE logical interfaces with an underlying Ethernet interface that already has a PPPoE dynamic profile attached.

```
[edit interfaces ge-1/0/1 unit 0 pppoe-underlying-options]
user@host# set dynamic-profile basic-pppoe-profile
```

5. (Optional) Enable duplicate protection to prevent activation of another dynamic PPPoE logical interface for the same client on the underlying interface.

```
[edit interfaces ge-1/0/1 unit 0 pppoe-underlying-options]
user@host# set duplicate-protection
```

6. (Optional) Configure the maximum number of concurrent PPPoE sessions that the router can activate on the underlying interface in either of the following ways:

   - To configure the maximum number of concurrent PPPoE sessions on a per-interface basis, from 1 to the platform-specific default for your router, use the **max-sessions** statement:

```
[edit interfaces ge-1/0/1 unit 0 pppoe-underlying-options]
user@host# set max-sessions 20
```

   - To configure the maximum number of concurrent PPPoE sessions on a per-subscriber basis, use the value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks vendor-specific attribute (VSA) [26-143]. By default, the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA takes precedence over the PPPoE maximum session value configured with the **max-sessions** statement.

7. (Optional) Configure the router to ignore the value returned by RADIUS in the Max-Clients-Per-Interface VSA and restore the PPPoE maximum session value on the underlying interface to the value configured in the CLI with the **max-sessions** statement.

```
[edit interfaces ge-1/0/1 unit 0 pppoe-underlying-options]
user@host# set max-sessions-vsa-ignore
```

8. (Optional) Enable PPPoE subscriber session lockout on the PPPoE underlying interface in either of the following ways:

   - To configure PPPoE subscriber session lockout with the default lockout period:

```
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options]
user@host# set short-cycle-protection
```

   - To configure PPPoE subscriber session lockout with a nondefault lockout period:

```
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options]
```

 Copyright © 2013, Juniper Networks, Inc.

```
user@host# set short-cycle-protection lockout-time-min minimum-seconds
  lockout-time-max maximum-seconds
```

BEST PRACTICE: When you configure PPPoE subscriber session lockout, we recommend that you also enable duplicate protection to ensure that the MAC source address for each PPPoE session is unique on the underlying interface.

9. (Optional) Specify the alternative name for the access concentrator, also known as the PPPoE server.

```
[edit interfaces ge-1/0/1 unit 0 pppoe-underlying-options]
user@host# set access-concentrator server-east
```

Related Documentation

- Subscriber Interfaces and PPPoE Overview on page 883
- Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview on page 887
- Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles on page 899
- Configuring the PPPoE Family for an Underlying Interface on page 831
- Configuring Lockout of PPPoE Subscriber Sessions on page 912
- Verifying and Managing Dynamic PPPoE Configuration on page 914
- Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Gigabit Ethernet VLAN Interface on page 917
- For information about configuring static Ethernet underlying interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface

You can limit the maximum number of concurrent static or dynamic PPPoE logical interfaces (sessions) that the router can activate on the PPPoE underlying interface, or the maximum number of active static or dynamic PPPoE sessions that the router can establish with a particular service entry in a PPPoE service name table.

You can configure the PPPoE maximum session limit in either of the following ways:

- On a per-interface basis, by using the **max-sessions** CLI statement.
- (Default) On a per-subscriber basis, by using the value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks vendor-specific attribute (VSA) [26-143]. By default, the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA takes precedence over the PPPoE maximum session value configured with the **max-sessions** statement.

Optionally, you can configure the router to ignore the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA by including the **max-sessions-vsa-ignore** statement in your configuration. Using the **max-sessions-vsa-ignore** statement restores the PPPoE maximum session value on the underlying interface to the value configured in the CLI with the **max-sessions** statement.

Before you begin:

- Configure the PPPoE underlying interface.

  To configure the underlying interface for use with a PPPoE dynamic profile, see "Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces" on page 905.

  To configure the PPPoE family for an underlying interface, see "Configuring the PPPoE Family for an Underlying Interface" on page 831.

To configure the PPPoE maximum session limit:

1. Specify that you want to configure PPPoE-specific options on the underlying interface:
   - For a PPPoE family in a dynamic profile for a VLAN demultiplexing (demux) logical interface:

     [edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number*]
     user@host# **edit family pppoe**

   - For a PPPoE family in a dynamic profile:

     [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]
     user@host# **edit family pppoe**

   - For a PPPoE underlying interface in a dynamic profile:

     [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]
     user@host# **edit pppoe-underlying-options**

   - For a PPPoE family on an underlying interface:

 Copyright © 2013, Juniper Networks, Inc.

[edit interfaces *interface-name* unit *logical-unit-number*]
user@host# **edit family pppoe**

- For an underlying interface with PPPoE encapsulation:

  [edit interfaces *interface-name* unit *logical-unit-number*]
  user@host# **edit** pppoe-underlying-options

- For an underlying interface established with a particular service entry in a PPPoE service name table:

  [edit protocols pppoe service-name-tables *table-name*]
  user@host# **edit service** ***service-name***

2. Configure the maximum number of concurrent PPPoE sessions that the router can activate on the underlying interface in either of the following ways:

   - To configure the maximum number of concurrent PPPoE sessions on a per-interface basis, from 1 to the platform-specific default for your router, use the **max-sessions** statement:

     [edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]
     user@host# **set** max-sessions ***number***

   - To configure the maximum number of concurrent PPPoE sessions on a per-subscriber basis, use the value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks vendor-specific attribute (VSA) [26-143]. By default, the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA takes precedence over the PPPoE maximum session value configured with the **max-sessions** statement.

3. (Optional) To restore the PPPoE maximum session value on the underlying interface to the value configured in the CLI with the **max-sessions** statement, configure the router to ignore the value returned by RADIUS in the Max-Clients-Per-Interface VSA.

   [edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]
   user@host# **set** max-sessions-vsa-ignore

   NOTE: **You can issue the max-sessions-vsa-ignore statement at the same hierarchy levels as the max-sessions statement, with the exception of the [edit protocols pppoe service-name-tables *table-name* service *service-name*] hierarchy level.**

Related Documentation

- PPPoE Maximum Session Limit Overview on page 890
- Guidelines for Using PPPoE Maximum Session Limit from RADIUS on page 892
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces on page 905
- Configuring the PPPoE Family for an Underlying Interface on page 831
- For more information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation

You can create a dynamic PPPoE subscriber interface based on the service name, agent circuit identifier (ACI), and agent remote identifier (ARI) information provided by the PPPoE client during PPPoE negotiation. To do so, you assign a previously configured PPPoE dynamic profile to a named service, **empty** service, or **any** service entry in a PPPoE service name table, or to an ACI/ARI pair defined for these services.

Similarly, to specify the routing instance in which to instantiate the dynamic PPPoE subscriber interface, you can assign a previously configured routing instance to a named service, **empty** service, or **any** service in a PPPoE service name table, or to an ACI/ARI pair defined for these services.

Observe the following configuration guidelines when you assign a dynamic profile and routing instance to a PPPoE service name table to create a dynamic PPPoE subscriber interface:

- The dynamic profile and routing instance must already be configured on the router.
- The dynamic profile or routing instance assigned to the PPPoE service name table overrides the dynamic profile or routing instance assigned to the PPPoE underlying interface on which the dynamic subscriber interface is created.
- You cannot configure a dynamic profile or routing instance for an ACI/ARI pair already configured with a static interface (by using the **static-interface** statement). Conversely, you cannot configure a static interface for an ACI/ARI pair already configured with a dynamic profile or routing instance.

Before you begin:

1. Configure a PPPoE dynamic profile in either of the following ways:
   - To configure a basic PPPoE dynamic profile, see "Configuring a Basic PPPoE Dynamic Profile" on page 900.
   - To configure a PPPoE dynamic profile with additional options for subscriber access, see "Configuring a PPPoE Dynamic Profile with Additional Options" on page 903.
2. Configure the routing instance in which you want the router to instantiate the dynamic profile.

   For information about configuring routing instances, see the *Junos OS Routing Protocols Library for Routing Devices*.
3. Create the PPPoE service name table on the router.

   See *Creating a Service Name Table* in the *Junos OS Network Interfaces Library for Routing Devices*.

Copyright © 2013, Juniper Networks, Inc.

To create a dynamic PPPoE subscriber interface based on the service name and, optionally, associated ACI/ARI pair configured in a PPPoE service name table, do one of the following:

- Assign a previously configured dynamic profile and routing instance to a named, **empty**, or **any** service.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service premium dynamic-profile premiumProfile routing-instance
    premiumRI
  ```

- Assign a previously configured dynamic profile and routing instance to the ACI/ARI pair defined for a named, **empty**, or **any** service.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service any agent-specifier aci west-ge-3/0/3 ari sunnyvale
    dynamic-profile standardProfile routing-instance standardRI
  ```

Related Documentation

- Example: Configuring a PPPoE Service Name Table for Dynamic Subscriber Interface Creation on page 919
- Subscriber Interfaces and PPPoE Overview on page 883
- Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles on page 899
- *Configuring PPPoE Service Name Tables*

Configuring Lockout of PPPoE Subscriber Sessions

You can configure the router to temporarily prevent (lock out) a failed or short-lived PPPoE subscriber session from reconnecting to the router for a default or configurable period of time. Configuring a lockout period on the PPPoE underlying interface for static or dynamic PPPoE subscriber sessions protects the router and any external authentication, authorization, and accounting (AAA) servers, such as RADIUS, from excessive loading as a result of failed or short-lived (also known as short-cycle) PPPoE subscriber sessions that occur repeatedly for the same subscriber.

You can configure the router to use the default PPPoE lockout period, 1 through 300 seconds (5 minutes). Alternatively, you can override the default lockout period by specifying a minimum lockout time and maximum lockout time, each of which can be from 1 through 86,400 seconds (24 hours).

Before you begin:

- Configure the PPPoE underlying interface.

  To configure the underlying interface for use with a PPPoE dynamic profile, see "Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces" on page 905.

  To configure the PPPoE family for an underlying interface, see "Configuring the PPPoE Family for an Underlying Interface" on page 831.

To configure temporary lockout of PPPoE subscriber sessions:

1. Specify that you want to configure PPPoE-specific options on the underlying interface:

   - For a PPPoE family in a dynamic profile for a VLAN demultiplexing (demux) logical interface:

     [edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number*]
     user@host# **edit family pppoe**

   - For a PPPoE family in a dynamic profile:

     [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]
     user@host# **edit family pppoe**

   - For a PPPoE underlying interface in a dynamic profile:

     [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]
     user@host# **edit pppoe-underlying-options**

   - For a PPPoE family on an underlying interface:

     [edit interfaces *interface-name* unit *logical-unit-number*]
     user@host# **edit family pppoe**

   - For an underlying interface with PPPoE encapsulation:

     [edit interfaces *interface-name* unit *logical-unit-number*]
     user@host# **edit pppoe-underlying-options**

 Copyright © 2013, Juniper Networks, Inc.

2. Enable duplicate protection to prevent negotiation of a dynamic or static PPPoE client session on the same underlying interface when a PPPoE client session with the same media access control (MAC) source address is already active on that interface.

   [edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]
   user@host# **set** duplicate-protection

   BEST PRACTICE: **When you configure PPPoE subscriber session lockout, we recommend that you enable duplicate protection to ensure that the MAC source address for each PPPoE session is unique on the underlying interface.**

3. Enable PPPoE subscriber session lockout on the PPPoE underlying interface in either of the following ways:

   - To configure PPPoE subscriber session lockout with the default lockout period:

     [edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]
     user@host# **set** short-cycle-protection

   - To configure PPPoE subscriber session lockout with a nondefault lockout period:

     [edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]
     user@host# **set** short-cycle-protection **lockout-time-min** ***minimum-seconds***
     lockout-time-max ***maximum-seconds***

Related Documentation

- PPPoE Subscriber Session Lockout Overview on page 893
- Clearing Lockout of PPPoE Subscriber Sessions on page 913
- Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces on page 905
- Configuring the PPPoE Family for an Underlying Interface on page 831
- For more information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

Clearing Lockout of PPPoE Subscriber Sessions

Purpose Clear the lockout condition for the PPPoE subscriber session associated with a unique MAC source address.

Action

- To clear the lockout condition for PPPoE subscriber sessions associated with all MAC source addresses on all underlying interfaces:

  user@host> **clear pppoe lockout**

- To clear the lockout condition for the PPPoE subscriber session associated with the specified MAC source address:

  user@host> **clear pppoe lockout mac-address** ***mac-address***

- To clear the lockout condition for all PPPoE subscriber sessions on the specified underlying interface:

user@host> **clear pppoe lockout underlying-interfaces** ***underlying-interface-name***

- To clear the lockout condition for the PPPoE subscriber session associated with the specified MAC source address on the specified underlying interface:

  user@host> **clear pppoe lockout mac-address** ***mac-address*** **underlying-interfaces** ***underlying-interface-name***

- To verify that the lockout condition has been cleared:

  user@host> **show pppoe lockout**

Related Documentation

- PPPoE Subscriber Session Lockout Overview on page 893
- Configuring Lockout of PPPoE Subscriber Sessions on page 912
- Verifying and Managing Dynamic PPPoE Configuration on page 914
- *Junos OS Operational Mode Commands*

Verifying and Managing Dynamic PPPoE Configuration

Purpose View or clear information about dynamic PPPoE logical interfaces, underlying interfaces for dynamic PPPoE logical interfaces, and PPPoE statistics.

Action

- To display information about the properties of all PPPoE underlying interfaces associated with a dynamic PPPoE profile:

  user@host> **show pppoe underlying-interfaces**

- To display information about the PPPoE properties of a specified underlying interface associated with a dynamic PPPoE profile:

  user@host> **show pppoe underlying-interfaces** ***interface-name***

- To display session-specific information about PPPoE interfaces, including whether the interface was dynamically created or statically created:

  user@host> **show pppoe interfaces**

- To display information for a specified PPPoE service name table, including the assigned dynamic profile and routing instance, if configured:

  user@ host> **show pppoe service-name-tables** ***table-name***

- To display information about all active PPPoE sessions on the router:

  user@host > **show pppoe sessions**

- To display information for all active PPPoE sessions established for a specified service name:

  user@host > **show pppoe sessions service** ***service-name***

- To display information for all active PPPoE sessions established for a specified agent circuit identifier (ACI) or agent remote identifier (ARI) string:

  user@host > **show pppoe sessions aci "west-ge-2/0/3"**
  user@host > **show pppoe sessions ari "sunnyvale"**

 Copyright © 2013, Juniper Networks, Inc.

- To display PPPoE control packet statistics for all PPPoE sessions:

  user@host> **show pppoe statistics**

- To display PPPoE control packet statistics for a specified PPPoE underlying interface:

  user@host> **show pppoe statistics** ***interface-name***

- To clear (reset) PPPoE control packet statistics for all PPPoE sessions:

  user@host> **clear pppoe statistics**

- To clear (reset) PPPoE control packet statistics for a specified underlying Ethernet interface:

  user@host> **clear pppoe statistics** ***underlying-interface-name***

- To display summary information about PPPoE subscriber sessions currently undergoing lockout or currently in a lockout grace period on all PPPoE underlying interfaces:

  user@host> **show pppoe lockout**

- To display summary information about PPPoE subscriber sessions currently undergoing lockout or currently in a lockout grace period on the specified PPPoE underlying interface:

  user@host> **show pppoe lockout** ***underlying-interface-name***

Related Documentation

- *Junos OS Operational Mode Commands*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 48

Dynamic PPPoE Subscriber Interfaces Examples

- Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Gigabit Ethernet VLAN Interface on page 917
- Example: Configuring a PPPoE Service Name Table for Dynamic Subscriber Interface Creation on page 919
- Evaluation Order for Matching Client Information in PPPoE Service Name Tables on page 922

Example: Configuring a Dynamic PPPoE Subscriber Interface on a Static Gigabit Ethernet VLAN Interface

This example shows how to configure a dynamic PPPoE subscriber interface on a statically configured Gigabit Ethernet VLAN underlying interface. When a PPPoE subscriber logs in on the underlying interface, the router creates the dynamic PPPoE subscriber interface with the attributes specified in the dynamic profile.

In this example, the dynamic PPPoE profile, **pppoe-profile-east**, defines options for PPPoE subscribers accessing the network, and includes the predefined dynamic variables **$junos-interface-unit**, which represents the logical unit number of the dynamic PPPoE logical interface, and **$junos-underlying-interface**, which represents the name of the underlying Ethernet interface. The **pppoe-profile-east** dynamic profile is assigned to the underlying Ethernet VLAN interface **ge-2/0/3.1** that is configured with PPPoE (**ppp-over-ether**) encapsulation.

When the router dynamically creates the PPPoE subscriber interface on **ge-2/0/3.1** in response to a subscriber login, the values of **$junos-interface-unit** and **$junos-underlying-interface** are dynamically replaced with the actual logical unit number and interface name, respectively, that are supplied by the network when the PPPoE subscriber logs in.

To configure a dynamic PPPoE subscriber interface:

1. Configure a dynamic profile to define the attributes of the dynamic PPPoE subscriber interface.

```
[edit]
dynamic-profiles {
```

```
pppoe-profile-east {
  interfaces {
    pp0 {
      unit "$junos-interface-unit" {
        ppp-options {
          chap;
        }
        pppoe-options {
          underlying-interface "$junos-underlying-interface";
          server;
        }
        keepalives interval 30;
        family inet {
          filter {
            input pppoe-input-filter-east;
            output pppoe-output-filter-east precedence 20;
          }
          service {
            input {
              service-set inputService-east;
              post-service-filter postService-east;
            }
            output {
              service-set outputService-east;
            }
          }
          address 6.6.6.1/32;
          unnumbered-address lo0.0;
        }
      }
    }
  }
}
```

2. Assign the dynamic PPPoE profile to the static underlying Ethernet interface, and define PPPoE-specific attributes for the underlying interface.

```
[edit]
interfaces {
  ge-2/0/3 {
    vlan-tagging;
    unit 1 {
      encapsulation ppp-over-ether;
      vlan-id 100;
      pppoe-underlying-options {
        access-concentrator server-east;
        duplicate-protection;
        dynamic-profile pppoe-profile-east;
        max-sessions 10;
      }
    }
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Subscriber Interfaces and PPPoE Overview on page 883
- Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview on page 887
- Configuring a PPPoE Dynamic Profile with Additional Options on page 903
- Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces on page 905

Example: Configuring a PPPoE Service Name Table for Dynamic Subscriber Interface Creation

This example shows how to configure a PPPoE service name table to create a dynamic PPPoE subscriber interface based on the service name, agent circuit identifier (ACI), and agent remote identifier (ARI) information provided by PPPoE clients during PPPoE negotiation.

In this example, PPPoE service name table **TableDynamicPPPoE** includes an **any** service entry, **empty** service entry, and two named service entries: **Premium** and **Standard**. The PPPoE underlying interfaces configured for **TableDynamicPPPoE** are **ge-2/0/0.1** and **ge-2/0/0.2**. Only **ge-2/0/0.1** is configured for dynamic profile assignment and creation of dynamic PPPoE subscriber interfaces.

Following the configuration example, Table 71 on page 921 explains how the router evaluates the entries in **TableDynamicPPPoE** to create a dynamic PPPoE subscriber interface in a specified routing instance for each of several sample clients.

To configure a PPPoE service name table to create dynamic PPPoE subscriber interfaces:

1. Configure the PPPoE service name table.

```
protocols {
  pppoe {
    service-name-tables TableDynamicPPPoE {
      service any {
        terminate;
        max-sessions 100;
        dynamic-profile AnyProfile;
        agent-specifier {
          aci "broadway-ge-1/0/1.0" ari "london" {
            terminate;
            dynamic-profile LondonProfile;
            routing-instance LondonRI;
          }
          aci "groton-ge-4/0/3.32" ari "paris" {
            delay 5;
            dynamic-profile ParisProfile;
            routing-instance ParisRI;
          }
        }
      }
      service empty {
        drop;
```

```
                agent-specifier {
                  aci "dunstable-ge-1/0/0.1" ari "kanata" {
                    dynamic-profile BasicPppoeProfile;
                    delay 10;
                  }
                }
              }
              service Premium {
                terminate;
                dynamic-profile PremiumProfile;
              }
              service Standard {
                terminate;
                max-sessions 10;
                dynamic-profile StandardProfile;
                agent-specifier {
                  aci "dunstable-ge-1/0/0.1" ari "kanata" {
                    dynamic-profile BasicPppoeProfile;
                    delay 10;
                  }
                }
              }
            }
          }
        }
```

2. Configure the PPPoE underlying interface for the service name table.

```
interfaces {
  ge-2/0/0 {
    vlan-tagging;
    unit 1 {
      vlan-id 1;
      pppoe-underlying-options {
        dynamic-profile BasicPppoeProfile;
        service-name-table TableDynamicPPPoE;
      }
    }
    unit 2 {
      vlan-id 2;
      pppoe-underlying-options {
        service-name-table TableDynamicPPPoE;
      }
    }
  }
}
```

Table 71 on page 921 lists the service name, ACI value, and ARI value provided in several sample PPPoE client requests, and the name of the PPPoE underlying interface on which the router received each client request. The Results column describes the dynamic PPPoE subscriber interface created by the router based on *both* of the following:

- The values received from each PPPoE client during PPPoE negotiation
- The sequence in which the router evaluates the entries configured in the PPPoE service name table to find a match for the client's service name and ACI/ARI information, as

 Copyright © 2013, Juniper Networks, Inc.

described in "Evaluation Order for Matching Client Information in PPPoE Service Name Tables" on page 922

Table 71: Dynamic PPPoE Subscriber Interface Creation Based on PPPoE Client Request Values

PPPoE Client	Service Name	ACI Value	ARI Value	Receiving Underlying Interface	Results
Client 1	Premium	broadway-ge-1/0/1.1	london	ge-2/0/0.1	Matches ACI/ARI pair configured for **any** service. Router creates dynamic PPPoE subscriber interface over **ge-2/0/0.1** using **LondonProfile** dynamic profile and **LondonRI** routing instance assigned to **any** service.
Client 2	Premium	dunstable-ge-1/0/1.0	toronto	ge-2/0/0.1	Matches base **Premium** service. Router creates dynamic PPPoE subscriber interface over **ge-2/0/0.1** using **PremiumProfile** dynamic profile and routing instance associated with **ge-2/0/0.1** underlying interface.
Client 3	empty	dunstable-ge-1/0/0.1	kanata	ge-2/0/0.1	Matches ACI/ARI pair configured for **empty** service and **Standard** service. Router creates dynamic PPPoE subscriber interface over **ge-2/0/0/.1** after a delay of 10 seconds. Router uses **BasicPPPoEProfile** dynamic profile and routing instance associated with **ge-2/0/0.1** underlying interface.
Client 4	empty	slinger-ge-1/0/0.1	chicago	ge-2/0/0.2	Because receiving underlying interface **ge-2/0/0.2** is *not* associated with a dynamic profile, router does not create a dynamic PPPoE subscriber interface, and drops any PADI or PADR control packets received from this client.
Client 5	Standard	slinger-ge-1/0/0.1	chicago	ge-2/0/0.1	Matches base **Standard** service. Router creates dynamic PPPoE subscriber interface over **ge-2/0/0.1** using **StandardProfile** dynamic profile and routing instance associated with **ge-2/0/0.1** underlying interface.

Related Documentation

- Evaluation Order for Matching Client Information in PPPoE Service Name Tables on page 922
- Subscriber Interfaces and PPPoE Overview on page 883
- *Understanding PPPoE Service Name Tables*

- *Configuring PPPoE Service Name Tables*

Evaluation Order for Matching Client Information in PPPoE Service Name Tables

When the router receives a service request from a PPPoE client, it evaluates the entries configured in the PPPoE service name table to find a match for the client's ACI/ARI information so it can take the appropriate action.

The order of evaluation is as follows:

1. The router evaluates the ACI/ARI information configured for the **any** service entry, and ignores the contents of the service name tag transmitted by the client.
2. If no match is found for the client information, the router evaluates the ACI/ARI information for the **empty** service entry and the named service entries. If an ACI/ARI pair is not configured for these service entries, the router evaluates the other attributes configured for the **empty** service and named services.
3. If there is still no match for the client information, the router evaluates the other attributes configured for the **any** service entry, and ignores both the ACI/ARI information for the **any** service and the contents of the service name tag transmitted by the client. If the **any** service is configured for the default action, **drop**, the router drops the PADR packet. If the **any** service is configured for a nondefault action (**terminate** or **delay**), the router evaluates the other attributes configured for the **any** service.

Related Documentation

- *Understanding PPPoE Service Name Tables*
- *Benefits of Configuring PPPoE Service Name Tables*
- *Configuring PPPoE Service Name Tables*
- Example: Configuring a PPPoE Service Name Table for Dynamic Subscriber Interface Creation on page 919
- *PPPoE Overview*
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 49

MPLS Pseudowire Subscriber Interfaces Overview

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 925
- CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 925
- CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 927
- CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 930

Pseudowire Subscriber Logical Interfaces Overview

Subscriber management supports the creation of subscriber interfaces over point-to-point MPLS pseudowires. The pseudowire subscriber interface capability enables service providers to extend an MPLS domain from the access-aggregation network to the service edge, where subscriber management is performed. Service providers can take advantage of MPLS capabilities such as failover, rerouting, and uniform MPLS label provisioning, while using a single pseudowire to service a large number of DHCP and PPPoE subscribers in the service network.

The pseudowire is a tunnel that is either an MPLS-based Layer 2 VPN or Layer 2 circuit. The pseudowire tunnel transports Ethernet encapsulated traffic from an access node (for example, a DSLAM or other aggregation device) to the MX Series router that hosts the subscriber management services. The termination of the pseudowire tunnel on the MX Series router is similar to a physical Ethernet termination, and is the point at which subscriber management functions are performed. A service provider can configure multiple pseudowires on a per-DSLAM basis and then provision support for a large number of subscribers on a specific pseudowire. Figure 17 on page 924 shows an MPLS network that provides subscriber management support.

At the access node end of the pseudowire, the subscriber traffic can be groomed into the pseudowire in a variety of ways, limited only by the number and types of interfaces that can be stacked on the pseudowire. You specify an anchor point, which identifies the logical tunnel interface that terminates the pseudowire tunnel at the access node.

Copyright © 2013, Juniper Networks, Inc.

Figure 17: MPLS Access Network with Subscriber Management Support

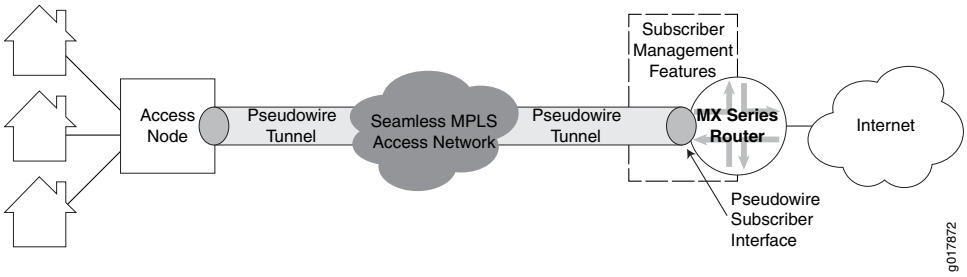

NOTE: Subscriber interfaces over MPLS pseudowires are supported on MX Series routers with MPCs.

Figure 18 on page 924 shows the protocol stack for a pseudowire subscriber logical interface. The pseudowire is a virtual device that is stacked above the logical tunnel anchor point on the physical interface (the IFD), and supports a circuit-oriented Layer 2 protocol (either Layer 2 VPN or Layer 2 circuit). The Layer 2 protocol provides the transport and service logical interfaces, and supports the protocol family (IPv4, IPv6, or PPPoE).

Figure 18: Pseudowire Subscriber Interface Protocol Stack

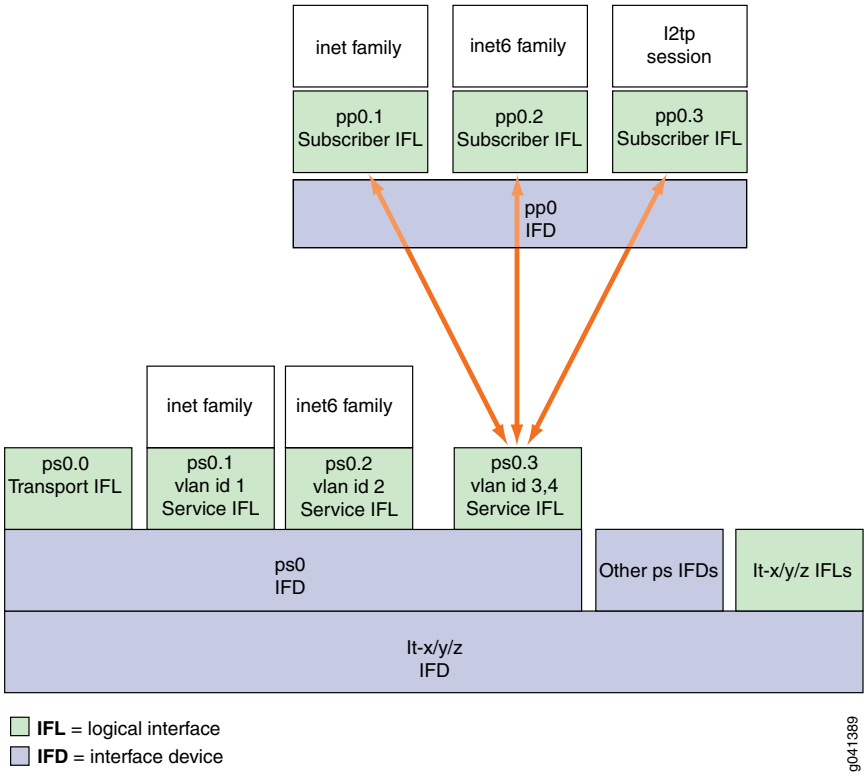

The pseudowire configuration is transparent to the subscriber management applications and has no impact on the packet payloads that are used for subscriber management.

Copyright © 2013, Juniper Networks, Inc.

Subscriber applications such as DHCP and PPPoE can be stacked over Layer 2 similar to the way in which they are stacked over a physical interface.

Related Documentation

- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 925

Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview

Junos OS supports two aspects of CoS for MPLS pseudowire subscriber interfaces. You can apply CoS rewrite rules and behavior aggregate (BA) classifiers to MPLS pseudowire subscriber interfaces. In addition, CoS performs egress hierarchical shaping towards the subscriber on MPLS pseudowire subscriber interfaces.

Hierarchical CoS enables you to apply traffic scheduling and queuing parameters and packet transmission scheduling parameters to an individual subscriber interface rather than to all interfaces configured on the port. Hierarchical CoS is supported on MX Series routers with either EQ DPCs or MPC/MICs installed.

On Juniper Networks MX Series routers, MPC/MIC and EQ DPC interfaces support a four-level CoS scheduling hierarchy that, when fully configured, consists of the physical interface (level 1), the interface set or the underlying interface (level 2), one or more logical interfaces (level 3), and one or more queues (level 4). Although all CoS scheduling hierarchies are four-level, level 1 is always the physical interface and level 4 is always the queue. Hierarchical scheduling configurations consist of the type of interfaces you configure; for example, a logical interface or an interface set and where those interfaces reside in the scheduling hierarchy, either level 2 or level 3. Because many hierarchical scheduling configurations are possible, we use the terms *two-level hierarchical scheduling* and *three-level hierarchical scheduling* in this discussion.

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 925
- CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 927
- CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 930
- **hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) on page 1677**

CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces

Two-level hierarchical scheduling limits the number of hierarchical levels in the scheduling hierarchy to two. In a two-level scheduling hierarchy, all logical interfaces and interface

sets share a single level 2 node. Table 72 on page 926 summarizes the interface hierarchy and the CoS scheduler node levels for two-level hierarchical scheduling.

Table 72: Two-Level Hierarchical Scheduling-Interface Hierarchy Versus Scheduling Nodes

Level 1	Level 2	Level 3	Level 4
Physical interface	–	Pseudowire transport logical interface	One or more queues
Physical interface	–	Interface set	One or more queues
Physical interface	–	Pseudowire service logical interface	One or more queues

You use the two-level hierarchical scheduling when you have many pseudowires but you do not require shaping specific to the subscriber logical interface. For example, when your configuration is one subscriber per pseudowire interface.

Figure 19 on page 926 shows a two-level hierarchical scheduling configuration for the MPLS pseudowires. In this configuration, level 1 is the physical interface used for the logical tunnel anchor node. All of the pseudowire transport interfaces share a single level 2 node. The level 3 nodes are the pseudowire transport logical interfaces (ps0.0, ps1.0, and ps2.0). In this configuration, interface sets are not configured and only the logical interfaces have traffic control profiles.

Figure 19: MPLS Pseudowire Subscriber Interface Two-Level Scheduler Configuration

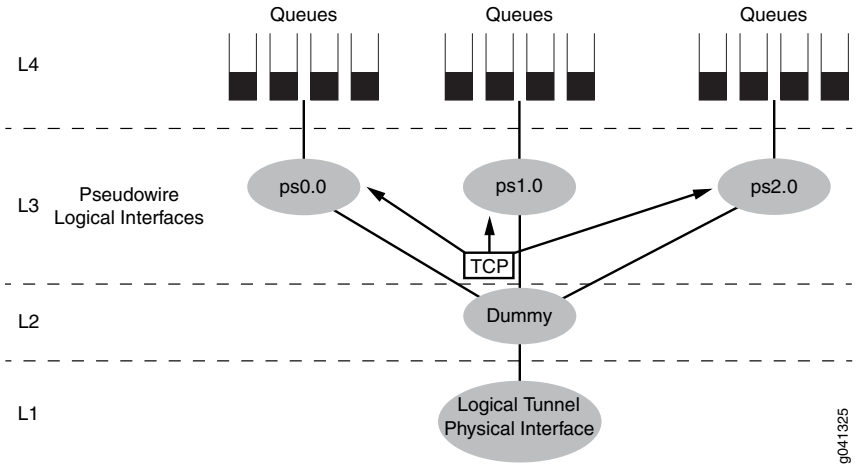

Two-level hierarchical scheduling has up to eight class of service queues. For this configuration, include the **maximum-hierarchy-levels 2** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement at the physical interface for the anchor logical tunnel.

Copyright © 2013, Juniper Networks, Inc.

NOTE: You cannot configure shaping policies on both the pseudowire logical interfaces and the subscriber logical interfaces over the same pseudowire. If a traffic-control profile is configured on a pseudowire logical interface, and CoS policies are configured on the subscriber logical interface over another pseudowire, all of the logical interfaces are at level 3 and act as peers.

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 925
- CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 927
- CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 930
- Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces on page 940
- **hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) on page 1677**

CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces

In three-level hierarchical scheduling, the CoS scheduler nodes at level 1, level 2, and level 3 form a scheduling hierarchy. You can configure many different three-level scheduling hierarchies, depending on the location of the interface set and the use of underlying interfaces. In all variations, the physical interface on which the logical tunnel resides is a level 1 CoS scheduler node and the queues reside at level 4. Three-level scheduling hierarchies can have up to eight class of service queues. Table 73 on page 927summarizes the most common three-level hierarchical scheduling configurations and shows the interface hierarchy and CoS scheduler nodes.

Table 73: Three-Level Hierarchical Scheduling-Interface Hierarchy Versus CoS Scheduling Node Levels

Level 1	Level 2	Level 3	Level 4
Physical interface	Pseudowire interface set	Pseudowire service logical interfaces	One or more queues
Physical interface	Pseudowire transport logical interface	Pseudowire interface set	One or more queues
Physical interface	Pseudowire transport logical interface	Pseudowire service logical interfaces	One or more queues

Three-Level Scheduling Hierarchy: Pseudowire Logical Interfaces over a Transport Logical Interface

Figure 20 on page 928 shows an MPLS pseudowire three-level scheduling hierarchy that includes two pseudowire service logical interfaces over a pseudowire transport logical interface. This variation uses the following scheduler nodes:

- Level 4—Forwarding class-based queues
- Level 3—Pseudowire service logical interfaces (ps0.1 and ps0.2) for subscriber sessions
- Level 2—Pseudowire transport logical interface (ps0.0)
- Level 1—Common/shared physical interface of the logical tunnel anchor point

You apply the traffic-control profiles at the pseudowire transport logical interfaces (level 2) and the pseudowire service logical interfaces (level 3).

Figure 20: Three-Level Scheduling Hierarchy Case 1: Pseudowire Service Logical Interfaces over a Transport Logical Interface

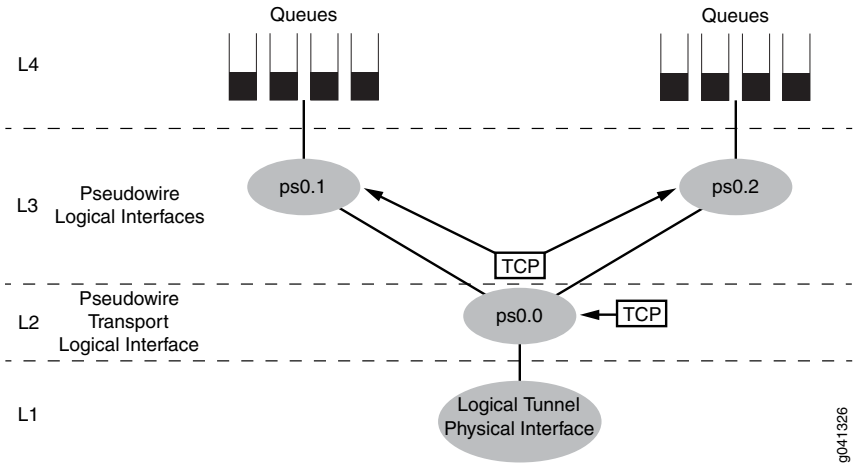

Three-Level Scheduling Hierarchy : Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set

Figure 21 on page 929 shows another variation of MPLS pseudowire three-level hierarchical scheduling that includes two pseudowire service logical interfaces over a pseudowire service interface set. This variation uses the following CoS scheduler nodes:

- Level 4—Forwarding class-based queues
- Level 3—Pseudowire service logical interfaces (ps0.1 and ps0.2)
- Level 2—Pseudowire service interface set
- Level 1—Common/shared physical interface of the logical tunnel anchor point

You apply the traffic-control profile at the pseudowire service interfaces (level 3) and at the interface set (level 2). This variation is most useful for subscriber edge deployments.

Copyright © 2013, Juniper Networks, Inc.

Figure 21: Three-Level Scheduling Hierarchy Case 2: Pseudowire Service Logical Interfaces over a Pseudowire Service Interface Set

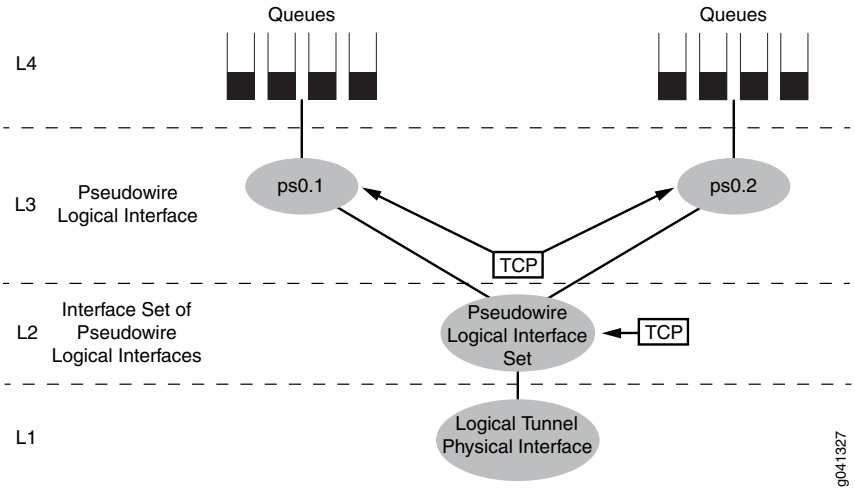

Three-Level Scheduling Hierarchy Combined Deployment Scenario

Figure 22 on page 929 shows a deployment scenario that combines the three-level hierarchical scheduling scenarios in Figure 20 on page 928 and Figure 21 on page 929.

Figure 22: Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces-Deployment Scenario

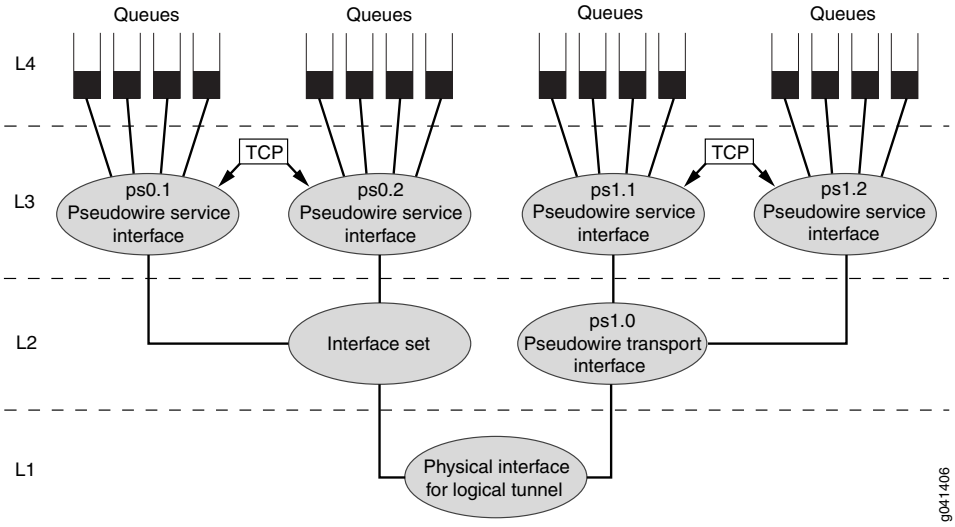

This variation uses the following CoS scheduler nodes:

- Level 4—Forwarding class-based queues
- Level 3—Pseudowire service logical interfaces (ps0.1, ps0.2, ps1.1, and ps1.2)

- Level 2—Service interface set for pseudowire service interfaces (ps0.1 and ps0.2) and transport logical interface (ps1.0) for the pseudowire service logical interfaces (ps1.1 and ps1.2)
- Level 1—Common/shared physical interface of the logical tunnel anchor point

You apply the traffic-control profiles to the interfaces at both level 2 and level 3, as well as the interface set at level 2.

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 925
- CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 930
- Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Transport Logical Interface) on page 942
- Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Pseudowire Interface Set) on page 944
- **hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) on page 1677**

CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces

CoS supports two-level and three-level hierarchies for MPLS pseudowire subscriber interfaces.

To configure two-level scheduling, include the **maximum-hierarchy-levels 2** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement on the physical interface of the logical tunnel anchor point.

To configure three-level hierarchical scheduling, include the **implicit-hierarchy** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement on the physical interface of the logical tunnel anchor point. Use the following guidelines for configuring the **implicit-hierarchy** option:

- If an output traffic-control profile is configured on the pseudowire transport interface and on a pseudowire service interface, the two interfaces form a scheduling hierarchy. The pseudowire transport interface resides in a level 2 scheduler node and the pseudowire service interface resides in a level 3 scheduler node.
- If an output traffic-control profile is configured on the pseudowire services interface but not on a pseudowire transport interface, the pseudowire services interface resides in a level 3 scheduler node.
- If an output traffic-control profile is only configured on the pseudowire transport interface and not on the pseudowire services interface, the pseudowire transport interface resides in a level 3 scheduler node and all pseudowire traffic uses this node.

 Copyright © 2013, Juniper Networks, Inc.

If the **implicit-hierarchy** option is not set on the logical tunnel anchor point, logical interfaces behave normally with the hierarchical-scheduler mode configured with or without the **hierarchical-scheduler maximum-hierarchy-levels** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement. In this case, when you apply a traffic-control profile to the pseudowire and service logical interfaces, they both reside in level 3 scheduler nodes and do not form a scheduling hierarchy, which might not be the desirable behavior. In business edge, where only the pseudowire logical interfaces need to be shaped, applying the traffic-control profile at just the transport logical interface may be sufficient.

When configuring the logical tunnel physical interface for the maximum hierarchy level, all pseudowire logical interfaces operating on the physical interface use the same hierarchy model. If you want to mix two-level and three-level scheduling hierarchies, you can group the pseudowires together by hierarchy levels and share the same logical tunnel anchor point or you can use three-level scheduling for all pseudowires over the anchor point.

To specify rewrite rules and classifiers on pseudowire interfaces, reference the pseudowire device under the **[edit class-of-service interfaces]** hierarchy level and specify the rewrite rules and classifiers for the pseudowire interfaces.

To control all pseudowire traffic using the same logical tunnel interface, apply CoS policies at the physical interface for the anchor logical tunnel.

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 925
- Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces on page 940
- Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Transport Logical Interface) on page 942
- Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Pseudowire Interface Set) on page 944
- **hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) on page 1677**

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 50

Configuring MPLS Pseudowire Subscriber Interfaces

- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router on page 935
- Configuring a Pseudowire Subscriber Logical Interface Device on page 935
- Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface on page 937
- Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces on page 937
- Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces on page 938
- Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface on page 939
- Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces on page 940
- Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Transport Logical Interface) on page 942
- Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Pseudowire Interface Set) on page 944

Configuring a Pseudowire Subscriber Logical Interface

A pseudowire subscriber logical interface terminates an MPLS pseudowire tunnel from an access node to the MX Series router that hosts subscriber management, and enables you to perform subscriber management services at the interface.

To create a pseudowire subscriber logical interface:

1. Specify the number of pseudowire logical interfaces that the router can support.

   See "Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router" on page 935.

2. Configure the pseudowire subscriber logical interface device.

Copyright © 2013, Juniper Networks, Inc.

See "Configuring a Pseudowire Subscriber Logical Interface Device" on page 935.

3. Configure the transport logical interface.

   See "Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface" on page 937.

4. Configure the signaling for the pseudowire subscriber interface. You can use either Layer 2 circuit signaling or Layer 2 VPN signaling. The two signaling types are mutually exclusive for a given pseudowire.

   - To configure Layer 2 circuit signaling, see "Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces" on page 937.
   - To configure Layer 2 VPN signaling, see "Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces" on page 938.

5. Configure the service logical interface.

   See "Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface" on page 939.

6. Configure the underlying interface device.

   See "Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces" on page 905.

7. Configure CoS parameters and BA classification.

   See "CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces" on page 930.

8. (Optional) Associate a dynamic profile with the pseudowire subscriber logical interface.

   You can associate DHCP, PPPoE, IP demux, and VLAN dynamic profiles with pseudowire subscriber logical interfaces. The support is similar to the typical Ethernet interface support.

   NOTE: **When using a PPPoE dynamic profile to create a pseudowire subscriber logical interface over a demux interface device, the dynamic profile must explicitly specify the correct pseudowire interface device over which the interface is created. The dynamic profile does not automatically create the interface over the demux0 interface device, as is the case with a VLAN demux interface.**

   For additional information about associating dynamic profiles to interfaces, see "Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview" on page 207 and "Configuring VLAN Interfaces to Use Dynamic Profiles" on page 720.

9. (Optional) Configure interface set support for pseudowire subscriber logical interfaces.

   See *Configuring Interface Sets* and *Applying Interface Sets*.

10. (Optional) Stack PPPoE logical interfaces over a pseudowire logical device.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923

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

You must set the maximum number of pseudowire logical interface devices (pseudowire tunnels) that the router can use for subscriber logical interfaces. You can specify a maximum of 2048 pseudowire logical interface devices for an MX Series router. Each pseudowire device supports a maximum of 4000 subscriber logical interfaces.

A PFE can host a maximum of 2048 pseudowire logical interface devices, which is the chassis maximum. This PFE hosting support provides the configuration flexibility needed for special cases that might occur for business edge scenarios. However, you can exceed the available PFE resources as you configure additional services on the pseudowire logical interface devices ports. To support a scaled configuration, ensure that you populate the appropriate number of PFEs for the chassis, and that you distribute the pseudowire logical interface devices across the PFEs in such a way that ensures that no PFE is overwhelmed by the anticipated peak load. As part of the network planning for your particular deployment, you must consider the exact mix of the distribution of the pseudowire logical interface devices and the services associated with the devices.

BEST PRACTICE: **A configured pseudowire logical interface device consumes resources from shared pools even when the device has no active subscriber logical interfaces. To conserve resources, do not deploy an excessive number of pseudowire devices that you do not intend to use.**

To configure the number of pseudowire logical interface devices that you want the router to support:

1. Specify that you want to configure the pseudowire service.

   ```
   [edit chassis]
   user@host# edit pseudowire-service
   ```

2. Set the maximum number of pseudowire logical interface devices.

   ```
   [edit chassis pseudowire-service]
   user@host# set device-count 500
   ```

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933

Configuring a Pseudowire Subscriber Logical Interface Device

To configure a pseudowire logical interface device that the router uses for subscriber logical interfaces, you specify the logical tunnel that processes the pseudowire

termination. You can also configure additional optional parameters for the interface device, such as VLAN tagging method, MTU, and gratuitous ARP support.

To configure the pseudowire subscriber interface device:

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

   user@host# **edit interfaces** ps0

2. Specify the logical tunnel interface that is the anchor point for the pseudowire logical interface device. The anchor point must be an **lt** device in the format **lt-*fpc/pic/port***.

   [edit interfaces ps0]
   user@host# **set** anchor-point **lt-1/0/10**

3. (Optional) Specify the VLAN tagging method used for the pseudowire logical interface device. You can specify single tagging, dual (stacked) tagging, mixed (flexible) tagging, or no tagging.

   [edit interfaces ps0]
   user@host# **set flexible-vlan-tagging**

   See *Enabling VLAN Tagging* for additional information about VLAN tagging.

4. (Optional) Specify the MTU for the pseudowire logical interface device. If you do not explicitly configure the MTU, the router uses the default value of 1500.

   [edit interfaces ps0]
   user@host# **set mtu 2500**

   See *Setting the Protocol MTU* for additional information.

5. (Optional) Specify that the pseudowire logical interface device does not respond to gratuitous ARP requests.

   [edit interfaces ps0]
   user@host# **set no-gratuitous-arp-request**

   See *Configuring Gratuitous ARP* for additional information.

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

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933

Copyright © 2013, Juniper Networks, Inc.

Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface

This topic describes how to configure a pseudowire transport logical interface. A pseudowire device can have only one transport logical interface.

A pseudowire logical device and its related pseudowire logical interfaces are dependent on the state of the underlying logical transport interface device, which is either the Layer 2 VPN or Layer 2 circuit.

NOTE: We recommend that you use unit 0 to represent the transport logical interface for the pseudowire device. Non-zero unit numbers represent *service* logical interfaces used for pseudowire subscriber interfaces.

To configure a pseudowire transport logical interface:

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

   ```
   [edit]
   user@host# edit interfaces ps0
   ```

2. Specify that you want to configure unit 0, which represents the transport logical interface.

   ```
   [edit interfaces ps0]
   user@host# edit unit 0
   ```

3. Specify the **ethernet-ccc** encapsulation method for the transport logical interface.

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

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933

Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces

This topic describes the steps for configuring Layer 2 circuit signaling used for the pseudowire subscriber logical interface support. You can also use Layer 2 VPN signaling for pseudowire subscriber logical interfaces. The two methods are mutually exclusive; you can use only one method for a particular pseudowire.

To configure Layer 2 circuit signaling for pseudowire interfaces:

1. Specify that you want to configure Layer 2 circuit parameters at the protocols hierarchy level.

   ```
   [edit protocols]
   user@host# edit l2circuit
   ```

2. Specify the IP address of the neighbor, to identify the PE router used for the Layer 2 circuit.

   [edit protocols l2circuit]
   user@host# **edit neighbor 192.168.102.15**

3. Specify the interface used by the Layer 2 circuit traffic.

   [edit protocols l2circuit neighbor 192.168.102.15]
   user@host# **edit interface ps1.0**

4. Configure the virtual circuit ID that identifies the Layer 2 circuit for the pseudowire.

   [edit protocols l2circuit neighbor 192.168.102.15 interface ps1.0]
   user@host# **set virtual-circuit-id 5**

5. (Optional) If multiple VLAN interfaces are carried over the pseudowire Layer 2 payload, configure the **no-vlan-id-validate** statement. This statement prevents VLAN validation in the signaling.

   [edit protocols l2circuit neighbor 192.168.102.15 interface ps1.0]
   user@host# **set** no-vlan-id-validate

For more information about Layer 2 circuits, see *Configuring Interfaces for Layer 2 Circuits*.

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933
- *Configuring Interfaces for Layer 2 Circuits*

Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces

This topic describes the steps for configuring Layer 2 VPN signaling used for the pseudowire subscriber logical interface support. You can also use Layer 2 circuit signaling for pseudowire subscriber logical interfaces. The two methods are mutually exclusive; you can use only one method on a particular pseudowire.

To configure Layer 2 VPN signaling for pseudowire interfaces:

1. Specify the name of the routing instance you want to configure.

   [edit]
   user@host# **edit routing-instances l2vpn0**

2. Configure the Layer 2 VPN routing instance type.

   [edit routing-instances l2vpn0]
   user@host# **set instance-type l2vpn**

3. Associate the pseudowire logical interface for the Layer 2 VPN.

   [edit routing-instances l2vpn0]
   user@host# **set interface ps1.0**

4. Configure the unique identifier for the routes that belong to the Layer 2 VPN.

   [edit routing-instances l2vpn0]
   user@host# **set route-distinguisher 111.1.1.1:100**

 Copyright © 2013, Juniper Networks, Inc.

5. Configure the VPN routing and forwarding (VRF) target of the routing instance.

   [edit routing-instances l2vpn0]
   user@host# **set vrf-target target:10:100**

6. Specify that you want to configure the Layer 2 VPN protocol for the routing instance.

   [edit routing-instances l2vpn0]
   user@host# **edit protocols l2vpn**

7. Configure the encapsulation type for the routing instance.

   [edit routing-instances l2vpn0 protocols l2vpn]
   user@host# **set encapsulation-type ethernet**

8. Specify the site name and site identifier for the Layer 2 VPN.

   [edit routing-instances l2vpn0 protocols l2vpn]
   user@host# **set site PE1 site-identifier 1**

9. Specify the interface that connects to the site, and the remote interface to which you want the specified interface to connect.

   [edit routing-instances l2vpn0 protocols l2vpn]
   user@host# **set interface ps1.0 remote-site-id 2**

10. Configure the tracing options for traffic that uses the Layer 2 VPN.

    [edit routing-instances l2vpn0 protocols l2vpn]
    user@host# **set traceoptions file l2vpn flag all**

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933

Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface

This topic describes how to configure a pseudowire service logical interface. Service logical interfaces represent the attachment circuits for pseudowire logical interfaces.

As described in the "Pseudowire Subscriber Logical Interfaces Overview" on page 923, you can choose whether to configure a service logical interface together with a higher subscriber logical interface, depending upon the business need. In a broadband edge configuration, the higher subscriber logical interface is the demarcation point for subscribers. However, in a business edge configuration, the service logical interface is the demarcation point for the business subscribers, and also serves as the subscriber logical interface, so no subscriber logical interfaces are explicitly configured.

NOTE: Non-zero unit numbers represent *service* logical interfaces used for pseudowire subscriber interfaces. Use unit 0 to represent the *transport* logical interface for the pseudowire device.

To configure a pseudowire service logical interface:

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

```
[edit]
user@host# edit interfaces ps0
```

2. Configure the unit for the service logical interface. Use a non-zero unit number.

```
[edit interfaces ps0]
user@host# edit unit 1
```

3. Configure the VLAN tag IDs.

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

4. Configure the interface to respond to ARP requests when the device has an active route to the ARP request target address.

```
[edit interfaces ps0 unit 1]
user@host# set proxy-arp
```

5. Specify that you want to configure the protocol family information. Pseudowire service logical interfaces support IPv4 (inet), IPv6 (inet6), and PPPoE (pppoe) protocol families.

   For example, to configure the IPv4 family:

   a. Specify that you want to configure IPv4.

```
[edit interfaces ps0 unit 1]
user@host# edit family inet
```

   b. Configure the parameters for the family.

```
[edit interfaces ps0 unit 1 family inet]
user@host# set filter input filter 1 output filter 4
user@host# set mac-validate loose
user@host# set input-hierarchical-policer policer-1
user@host# set unnumbered-address lo0.0 preferred-source-address 100.0.0.1
```

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933

Configuring CoS Two-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces

Before configuring CoS parameters for MPLS pseudowire subscriber interfaces, you must first complete these tasks:

1. Configure the pseudowire logical interfaces. See "Configuring a Pseudowire Subscriber Logical Interface" on page 933.
2. Configure the pseudowire device count. See "Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router" on page 935.
3. Configure the pseudowire device including the logical tunnel anchor point. See "Configuring a Pseudowire Subscriber Logical Interface Device" on page 935.

Copyright © 2013, Juniper Networks, Inc.

4. Configure the pseudowire transport logical interface. See "Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface" on page 937.
5. Configure the pseudowire signaling (either Layer 2 circuit signaling or Layer 2 VPN signaling). See "Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces" on page 937 or "Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces" on page 938.
6. Configure the pseudowire logical interfaces. See "Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface" on page 939.

To configure CoS policies on MPLS pseudowire subscriber interfaces using two-level scheduling:

1. Configure the hierarchical scheduler for the physical interface used for the logical tunnel (anchor point). For two-level scheduling the hierarchical scheduler must be set to **maximum-scheduler levels 2**.

   [edit]
   user@host#**edit interfaces ps** *ps-anchor-device-name*
   user@host#**set hierarchical-scheduler maximum-hierarchy-levels 2**

2. Specify the traffic-control profile to use on the pseudowire logical interface.

   [edit class-of-service]
   user@host#**edit interfaces ps** *ps-device-name*
   user@host#**edit unit** *logical-unit-number*
   user@host#**set output-traffic-control-profile** *profile-name*

3. Configure the rewrite rule.

   The available rewrite rule types for pseudowire interfaces are **dscp** and **inet-precedence**.

   [edit class-of-service]
   user@host#**edit interfaces ps** *ps-device-name*
   user@host#**edit unit** *logical-unit-number*
   user@host#**edit rewrite-rules (dscp | inet-precedence)** *rewrite-name*
   user@host#**edit forwarding-class** *class-name*
   user@host#**set loss-priority** *class-name* **code-point (***alias* **|** *bits***)**

4. Configure the classifier.

   The available classifier types for pseudowire interfaces are **dscp** and **inet-precedence**.

   [edit class-of-service]
   user@host#**edit interfaces ps** *ps-device-name*
   user@host#**edit unit** *logical-unit-number*
   user@host#**edit classifiers (dscp | inet-precedence)** *classifier-name*
   user@host#**edit forwarding-class** *class-name*
   user@host#**set loss-priority** *class-name* **code-points [***aliases***] [***bit-patterns***]**

5. Apply the rewrite rule and classifier to the pseudowire interface.

   For the *interface_name* parameter, specify the pseudowire device name.

   [edit class-of-service interfaces *interface_name* unit *logical-unit-number*]
   user@host#**set rewrite-rule (dscp | inet-precedence) (***rewrite-name* **| default) protocol** *protocol-types*
   user@host#**set classifiers (dscp | inet-precedence) (***classifier-name* **| default)**

Related Documentation

- *CoS on Ethernet Pseudowires in Universal Edge Networks Overview*
- For more information about rewrite rules and classifiers, see the *Junos OS Class of Service Library for Routing Devices*
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 925
- CoS Two-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 925
- CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 927
- CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 930
- **hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) on page 1677**

Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Transport Logical Interface)

Before configuring CoS three-level scheduling on pseudowire logical interfaces over a transport logical interface, you must first complete these tasks:

1. Configure the pseudowire logical interfaces. See "Configuring a Pseudowire Subscriber Logical Interface" on page 933.
2. Configure the pseudowire device count. See "Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router" on page 935.
3. Configure the pseudowire device including the logical tunnel anchor point. See "Configuring a Pseudowire Subscriber Logical Interface Device" on page 935.
4. Configure the pseudowire transport logical interface. See "Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface" on page 937.
5. Configure the pseudowire signaling (either Layer 2 circuit signaling or Layer 2 VPN signaling). See "Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces" on page 937 or "Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces" on page 938.
6. Configure the pseudowire logical interfaces. See "Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface" on page 939.

Three-level scheduling on pseudowire logical interfaces over a transport logical interface requires you to apply the traffic-control profiles at both the pseudowire logical interface and the pseudowire transport logical interface. To configure CoS policies on three-level scheduling on pseudowire logical interfaces over a transport logical interface:

1. Configure the hierarchical scheduler for the physical interface used for the logical tunnel (anchor point). For three-level scheduling the hierarchical scheduler must be set to **implicit-hierarchy**.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit]
user@host#edit interfaces ps-anchor-device-name
user@host#set hierarchical-scheduler implicit-hierarchy
```

2. Specify the traffic-control profile to use on the pseudowire logical interface.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#set output-traffic-control-profile profile-name
```

3. Specify the traffic-control profile to use on the pseudowire transport logical interface.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#set output-traffic-control-profile profile-name
```

4. Configure the rewrite rule.

   The available rewrite rule types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit rewrite-rules (dscp | inet-precedence) rewrite-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-point (alias | bits)
```

5. Configure the classifier.

   The available classifier types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit classifiers (dscp | inet-precedence) classifier-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-points [aliases] [bit-patterns]
```

6. Apply the rewrite rule and classifier to the pseudowire interfaces.

   For the *interface_name* parameter, specify the pseudowire device name.

```
[edit class-of-service interfaces interface_name unit logical-unit-number]
user@host#set rewrite-rule (dscp | inet-precedence) (rewrite-name | default) protocol
  protocol-types
user@host#set classifiers (dscp | inet-precedence) (classifier-name | default)
```

Related Documentation

- *CoS on Ethernet Pseudowires in Universal Edge Networks Overview*
- For more information about rewrite rules and classifiers, see the *Junos OS Class of Service Library for Routing Devices*
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 925
- CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 927

- CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 930
- Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Pseudowire Interface Set) on page 944
- **hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) on page 1677**

Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Pseudowire Interface Set)

Before configuring three-level scheduling on pseudowire logical interfaces over a pseudowire logical interface set, you must first complete the following tasks:

1. Configure the pseudowire logical interfaces. See "Configuring a Pseudowire Subscriber Logical Interface" on page 933.
2. Configure the pseudowire device count. See "Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router" on page 935.
3. Configure the pseudowire device including the logical tunnel anchor point. See "Configuring a Pseudowire Subscriber Logical Interface Device" on page 935.
4. Configure the pseudowire transport logical interface. See "Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface" on page 937.
5. Configure the pseudowire signaling (either Layer 2 circuit signaling or Layer 2 VPN signaling). See "Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces" on page 937 or "Configuring Layer 2 VPN Signaling for Pseudowire Subscriber Logical Interfaces" on page 938.
6. Configure the pseudowire logical interfaces. See "Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface" on page 939.

Three-level scheduling on pseudowire logical interfaces over a pseudowire logical interface set requires you to apply the traffic-control profiles at both the pseudowire logical interface and the pseudowire logical interface-set. To configure CoS policies on MPLS pseudowire subscriber interfaces using three-level implicit hierarchical scheduling:

1. Configure the hierarchical scheduler for the physical interface used for the logical tunnel (anchor point). For three-level scheduling the hierarchical scheduler must be set to **implicit-hierarchy**.

   [edit]
   user@host#**edit interfaces** ***ps-anchor-device-name***
   user@host#**set hierarchical-scheduler implicit-hierarchy**

2. Specify the traffic-control profile to use on the pseudowire logical interfaces.

   [edit class-of-service]
   user@host#**edit interfaces ps** ***ps-device-name***
   user@host#**edit unit** ***logical-unit-number***
   user@host#**set output-traffic-control-profile** ***profile-name***

3. Define a pseudowire logical interface set and configure the traffic-control profile used for the interface set.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit class-of-service]
user@host#edit interfaces
user@host#edit interface-set interface-set-name
user@host#edit output-traffic-control-profile profile-name
```

4. Group the pseudowire logical interfaces in the pseudowire logical interface set.

```
[edit ]
user@host#edit interfaces
user@host#edit interface-set interface-set-name
user@host#edit interface ps ps-device-name
user@host#edit unit logical-unit-number
```

5. Configure the rewrite rule.

   The available rewrite rule types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit rewrite-rules (dscp | inet-precedence) rewrite-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-point (alias | bits)
```

6. Configure the classifier.

   The available classifier types for pseudowire interfaces are **dscp** and **inet-precedence**.

```
[edit class-of-service]
user@host#edit interfaces ps ps-device-name
user@host#edit unit logical-unit-number
user@host#edit classifiers (dscp | inet-precedence) classifier-name
user@host#edit forwarding-class class-name
user@host#set loss-priority class-name code-points [aliases] [bit-patterns]
```

7. Apply the rewrite rule and classifier to the pseudowire interfaces.

   For the *interface_name* parameter, specify the ps device name.

```
[edit class-of-service interfaces interface_name unit logical-unit-number]
user@host#set rewrite-rule (dscp | inet-precedence) (rewrite-name | default) protocol
  protocol-types
user@host#set classifiers (dscp | inet-precedence) (classifier-name | default)
```

Related Documentation

- *CoS on Ethernet Pseudowires in Universal Edge Networks Overview*
- For more information about rewrite rules and classifiers, see the *Junos OS Class of Service Library for Routing Devices*
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 925
- CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 927
- CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces on page 930

- Configuring CoS Three-Level Hierarchical Scheduling for MPLS Pseudowire Subscriber Interfaces (Logical Interfaces over a Transport Logical Interface) on page 942
- **hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) on page 1677**

Copyright © 2013, Juniper Networks, Inc.

PART 13

Class of Service for Subscriber Access

- Dynamic CoS for Subscriber Access Overview on page 949
- Configuration Summary of Dynamic CoS for Subscriber Access on page 963
- Configuring Dynamic Shaping and Scheduling for Subscriber Access on page 971
- RADIUS and Dynamic CoS Overview on page 989
- Configuring RADIUS for Dynamic CoS on page 999
- Interface Solutions for Dynamic CoS Overview on page 1013
- Configuring Interface Solutions for Dynamic CoS on page 1023
- Dynamic CoS for Subscriber Access Examples on page 1031
- Bandwidth Management for Dynamic CoS Overview on page 1069
- Configuring Bandwidth Management Parameters for Dynamic CoS on page 1089
- Bandwidth Management for Dynamic CoS Examples on page 1111

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 51

Dynamic CoS for Subscriber Access Overview

- CoS for Subscriber Access Overview on page 949
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956

CoS for Subscriber Access Overview

This topic describes class-of-service (CoS) functionality for dynamic subscriber access.

Junos CoS enables you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. This functionality allows packet loss to happen according to rules that you configure. The Junos CoS features provide a set of mechanisms that you can use to provide differentiated services when best-effort traffic delivery is insufficient.

In a subscriber access environment, service providers want to provide video, voice, and data services over the same network for subscribers. Subscriber traffic is delivered from the access network, through a router, through a switched Ethernet network, to an Ethernet digital subscriber line access multiplexer (DSLAM). The DSLAM forwards the subscriber's traffic to the residential gateway over a digital subscriber line (DSL). An MX Series router that is installed in a subscriber access network as an edge router can perform subscriber management functions that include subscriber identification and per-subscriber CoS.

In a subscriber access network, a subscriber is an authenticated user—a user that has logged in to the access network at a subscriber interface and then been verified by the configured authentication server and subsequently granted initial CoS services. Subscribers can be identified statically or dynamically. In this network, subscribers are mapped to VLANs, demux, or PPPoE interfaces.

You can configure the router to provide *hierarchical scheduling* or *per-unit scheduling* for subscribers:

- Hierarchical CoS enables you to apply traffic scheduling and queuing parameters (which can include a delay-buffer bandwidth) and packet transmission scheduling parameters (which can include buffer management parameters) to an individual

Copyright © 2013, Juniper Networks, Inc.

subscriber interface rather than to all interfaces configured on the port. Hierarchical CoS enables you to dynamically modify queues when subscribers require services.

- Per-unit scheduling enables one set of output queues for each logical interface configured under the physical interface. In per-unit scheduling configurations, each Layer 3 scheduler node is allocated a dedicated set of queues.

Because the interface sets corresponding to VLANs using agent-circuit-identifier information are created dynamically, you can apply CoS attributes, such as shaping, at the household level. You must set and define the CoS policy for the agent-circuit-identifier virtual VLAN interface set using the dynamic profile for the agent-circuit-identifier interface set (not the subscriber profile). CoS on dynamic VLANs includes support for level 3 or level 2 scheduler nodes for a dynamic interface set. You can also configure a traffic control profile and a remaining traffic control profile for a dynamic interface set. CoS on dynamic VLANs enables you to configure a dynamic scheduler map for a traffic control profile that is used by a dynamic interface set. In this case, the dynamic scheduler map must use the unique ID format.

Related Documentation

- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 963
- Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 965
- Configuring Per-Unit Scheduling in a Dynamic Profile for Subscriber Access on page 967

Guidelines for Configuring Dynamic CoS for Subscriber Access

This topic describes the hardware requirements and guidelines for configuring dynamic CoS in a subscriber access environment.

Hardware Requirements for Dynamic CoS

Table 74 on page 951 lists the hardware requirements based on subscriber interface type for the hierarchical scheduling and per-unit scheduling dynamic CoS configurations.

 Copyright © 2013, Juniper Networks, Inc.

Table 74: Hardware Required for Dynamic CoS Configurations

Dynamic CoS Configuration	Subscriber Interface Type	EQ DPCs on MX Series Routers	MPC/MIC Modules on MX Series Routers	IQ2 PICs on M120 and M320 Routers	IQ2E PICs on M120 and M320 Routers
Hierarchical CoS	Static and dynamic VLANs	Yes	Yes	No	No
	Static and dynamic VLANs over aggregated Ethernet	Yes	Yes	No	No
	Static or dynamic IP demux interfaces	Yes	Yes	No	No
	Static or dynamic IP demux interfaces over aggregated Ethernet	Yes	Yes	No	No
	Static or dynamic VLAN demux interfaces	No	Yes	No	No
	Static or dynamic VLAN demux interfaces over aggregated Ethernet	No	Yes	No	No
	Static PPPoE interfaces	No	Yes	Yes	Yes
	Dynamic PPPoE interfaces	No	Yes	No	Yes
	Static or dynamic PPPoE interfaces over aggregated Ethernet	No	Yes	No	No
	L2TP LAC tunnel over PPP	No	Yes	No	No
	L2TP LNS inline service over PPP	No	Yes	No	No

Copyright © 2013, Juniper Networks, Inc.

Table 74: Hardware Required for Dynamic CoS Configurations *(continued)*

Dynamic CoS Configuration	Subscriber Interface Type	EQ DPCs on MX Series Routers	MPC/MIC Modules on MX Series Routers	IQ2 PICs on M120 and M320 Routers	IQ2E PICs on M120 and M320 Routers
Per-unit scheduling	Static and dynamic VLANs	Yes	Yes	No	No
	Static and dynamic VLANs over aggregated Ethernet	No	No	No	No
	Static or dynamic IP demux interfaces	Yes	No	No	No
	Static or dynamic IP demux interfaces over aggregated Ethernet	No	No	No	No
	Static or dynamic VLAN demux interfaces	No	No	No	No
	Static or dynamic VLAN demux interfaces over aggregated Ethernet	No	No	No	No
	Static PPPoE interfaces	No	Yes	Yes	Yes
	Dynamic PPPoE interfaces	No	No	Yes	Yes
	Static or dynamic PPPoE interfaces over aggregated Ethernet	No	No	No	No
	L2TP LAC tunnel over PPP	No	No	No	No
	L2TP LNS inline service over PPP	No	No	No	No

Copyright © 2013, Juniper Networks, Inc.

Configuration Guidelines for Dynamic Scheduling and Queuing

When configuring scheduling and queuing for subscriber access, consider the following guidelines:

- To improve CoS performance in IPv4, IPv6, and dual-stack networks that use a DHCP access model, we recommend that you use the **aggregate-clients replace** statement rather than the **aggregate-clients merge** statement.
- You can configure dynamic CoS with one of the following scheduling configurations:
  - For hierarchical scheduling configurations, you must enable hierarchical scheduling in the static CLI for the interface referenced in the dynamic profile. If not, the dynamic profile fails.
  - For per-unit scheduling configurations, you must enable per-unit scheduling in the static CLI for the interface referenced in the dynamic profile. If not, the dynamic profile fails and schedulers are not attached to the interface.
- You configure the traffic scheduling and shaping parameters in a traffic-control profile within the dynamic profile. You can configure the scheduler map and schedulers in a dynamic profile or in the **[edit class-of-service]** hierarchy. You must statically configure the remaining CoS parameters, such as hierarchical scheduling, classifiers, drop profiles, and forwarding classes, in the **[edit class-of-service]** hierarchy.
- You can configure only one traffic-control-profile under a dynamic profile.
- You must define the output-traffic-control-profile that binds the traffic-control profile to the interface within the same dynamic profile as the interface.
- We recommend that you provide different names for the schedulers defined in dynamic profiles that are used for access and services. For example, if there are two dynamic profiles, voice-profile and video-profile, provide unique names for the schedulers defined under those profiles.
- You must use a service dynamic profile with a different profile name for each RADIUS CoA request over the same logical interface.
- When you configure scheduler and scheduler map sharing in client profiles, schedulers and scheduler maps must use the unique ID format. If the client profile uses the unique ID format and you want to have either scheduler or scheduler map sharing for service activation, you must configure the service profile in unique ID format.

Configuration Guidelines for Dynamic Classifiers and Rewrite Rules

When you configure classifiers and rewrite rules for subscriber access, consider the following guidelines:

- To apply classifiers and rewrite rules to a subscriber interface in a dynamic profile, you must configure the rewrite rule and classifier definitions in the static **[edit class-of-service]** hierarchy and reference them in the dynamic profile.

  - If a static classifier or a rewrite rule definition that is referenced by a dynamic subscriber interface does not exist, the configuration is invalid and the subscriber cannot log in.
  - If a network administrator changes the static classifiers and rewrite rules definitions that are referenced in a dynamic profile with an active subscriber interface logged in, the changes are applied to the active subscriber interface immediately.
  - If a network administrator deletes a classifier or a rewrite rule definition that is referenced by an active dynamic subscriber interface, the system removes the classifier or rewrite rule binding from the interface. The classifier is replaced by the default classifier. If the network administrator adds the removed classifier or rewrite rule to the configuration while the dynamic interface is active, the addition does not take effect until the subscriber logs out and then logs in again.

- IP demux interfaces can only instantiate Layer 3 rules (both rewrite rules and classifiers).
  - An IP demux subscriber interface can implicitly inherit a classifier from the underlying interface. If an IP demux interface is created without a classifier and a Layer 2 classifier is attached to the underlying interface, the IP demux interface also inherits the Layer 2 classifier. The **show class-of-service interface *interface-name*** command does not display this attachment.

    Table 75 on page 954 lists the classification rule configuration for an IP demux subscriber interface with a VLAN underlying interface.

Table 75: IP Demux Classification Rules

VLAN Underlying Interface Classifier Configuration	IP Demux Interface Classifier Configuration	Resulting Classifier Configuration
Layer 2	—	VLAN Layer 2
Layer 2	Layer 3	Demux Layer 3
Layer 3	—	Default
Layer 3	Layer 3	Demux Layer 3

  - An IP demux subscriber interface explicitly inherits Layer 2 rewrite rules from the underlying interface if a Layer 2 rewrite rule is present. The **show class-of-service interface *interface-name*** command displays the attachment.

    Table 76 on page 954 lists the rewrite rule configuration for an IP demux subscriber interface with a VLAN underlying interface.

Table 76: IP Demux Rewrite Rules

VLAN Underlying Interface Rewrite Rule Configuration	IP Demux Interface Rewrite Rule Configuration	Resulting Rewrite Rule Configuration
Layer 2	—	VLAN Layer 2

Copyright © 2013, Juniper Networks, Inc.

Table 76: IP Demux Rewrite Rules *(continued)*

VLAN Underlying Interface Rewrite Rule Configuration	IP Demux Interface Rewrite Rule Configuration	Resulting Rewrite Rule Configuration
Layer 2	Layer 3	VLAN Layer 2 and demux Layer 3
Layer 3	—	Default
Layer 3	Layer 3	Demux Layer 3

- An L2TP subscriber interface can implicitly inherit a classifier from the underlying interface.

  Table 77 on page 955 lists the classification rule configuration for an L2TP LAC subscriber interface with a VLAN underlying interface.

Table 77: L2TP Classification Rules

VLAN Underlying Interface Classifier Configuration	L2TP LAC Classifier Configuration	Resulting Classifier Configuration
Layer 2 or Fixed	Layer 2 or Fixed	VLAN Layer 2 or Fixed
Layer 2 or Fixed	Layer 3	Demux/PPPoE Layer 3
Layer 3	Layer 2 or Fixed	VLAN Layer 2 or Fixed
Layer 3	Layer 3	Demux/PPPoE Layer 3

- An L2TP LAC subscriber interface explicitly inherits Layer 2 rewrite rules from the underlying interface if a Layer 2 rewrite rule is present. Table 78 on page 955 lists the rewrite rule configuration for an L2TP LAC subscriber interface with a VLAN underlying interface.

Table 78: L2TP LAC Rewrite Rules

VLAN Underlying Interface Rewrite Rule Configuration	L2TP Interface Rewrite Rule Configuration	Resulting Rewrite Rule Configuration
Layer 2	Layer 2	VLAN Layer 2
Layer 2	Layer 3	VLAN Layer 2 and demux/PPPoE Layer 3
Layer 3	Layer 2	VLAN Layer 2 and demux/PPPoE Layer 3
Layer 3	Layer 3	Demux/PPPoE Layer 3

Copyright © 2013, Juniper Networks, Inc.

Configuration Guidelines for Dynamic CoS on Specific Interface Types

To obtain configuration guidelines for CoS on specific interface types, see:

- CoS for Aggregated Ethernet Subscriber Interfaces Overview on page 1013
- CoS for L2TP LAC Subscriber Interfaces Overview on page 1016
- CoS for L2TP LNS Inline Services Overview on page 1018
- CoS for PPPoE Subscriber Interfaces Overview on page 1014
- CoS for Interface Sets of Subscribers Overview on page 1019

Related Documentation

- CoS for Subscriber Access Overview on page 949
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 963
- Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 965
- Configuring Per-Unit Scheduling in a Dynamic Profile for Subscriber Access on page 967
- For information about static CoS configurations, see the *Junos OS Class of Service Library for Routing Devices*

Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers

This topic describes hierarchical CoS on MX Series routers running Modular Port Concentrator/Modular Interface Card (MPC/MIC) or Enhanced Queuing Dense Port Concentrator (EQ DPC) interfaces.

Hierarchical CoS enables you to apply traffic scheduling and queuing parameters and packet transmission scheduling parameters to an individual subscriber interface rather than to all interfaces configured on the port. Hierarchical CoS is supported on MX Series routers with either EQ DPCs or MPC/MICs installed.

On Juniper Networks MX Series routers, MPC/MIC and EQ DPC interfaces support a four-level CoS scheduling hierarchy that, when fully configured, consists of the physical interface (level 1), an interface set or underlying interface (level 2), one or more logical interfaces (level 3), and one or more queues (level 4). Although all CoS scheduling hierarchies are four-level, level 1 is always the physical interface and level 4 is always the queue. Hierarchical scheduling configurations consist of the type of interfaces you configure; for example, a logical interface or an interface set and where those interfaces reside in the scheduling hierarchy, either level 2 or level 3. Because many hierarchical scheduling configurations are possible, we use the terms *two-level hierarchical scheduling* and *three-level hierarchical scheduling* in this discussion.

 Copyright © 2013, Juniper Networks, Inc.

Two-Level Hierarchical Scheduling

Two-level hierarchical scheduling limits the number of hierarchical levels in the scheduling hierarchy to two (level 2 and level 3) as shown in Figure 23 on page 957. In this configuration, interface sets are not configured and only the logical interfaces have traffic control profiles. Two-level hierarchical scheduling is supported on MX Series routers running either MPC/MIC or EQ DPC interfaces.

Figure 23: Two-Level Hierarchical Scheduling

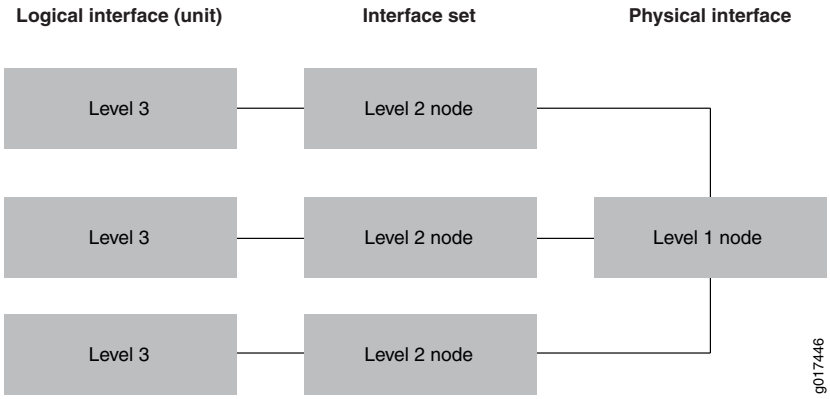

In a two-level scheduling hierarchy, all logical interfaces and interface sets share a single level 2 node; no hierarchical relationship is formed.

You control two-level hierarchical scheduling by including the **maximum-hierarchy-levels** option under the **[edit interfaces *interface-name* hierarchical-scheduler]** statement:

- When the **maximum-hierarchy-levels** option is not set, interface sets can be at either level 2 or level 3, depending on whether the member logical interfaces within the interface set have a traffic control profile.
- If any member logical interface has a traffic control profile, then the interface set is always a level 2 CoS scheduler node.
- If no member logical interface has a traffic control profile, the interface set is always a level 3 CoS scheduler node.
- If the **maximum-hierarchy-levels** option is set, then the interface set can only be at level 3; it cannot be at level 2. In this case, if you configure a level 2 interface set, you generate Packet Forwarding Engine errors.

Table 79 on page 957 summarizes the interface hierarchy and the CoS scheduler node levels for two-level hierarchical scheduling.

Table 79: Two-Level Hierarchical Scheduling-Interface Hierarchy Versus Scheduling Nodes

Level 1	Level 2	Level 3	Level 4
Physical interface	–	Logical interfaces	One or more queues

Copyright © 2013, Juniper Networks, Inc.

Table 79: Two-Level Hierarchical Scheduling-Interface Hierarchy Versus Scheduling Nodes *(continued)*

Level 1	Level 2	Level 3	Level 4
Physical interface	–	Interface set	One or more queues
Physical interface	–	Logical interfaces	One or more queues

To configure two-level hierarchical scheduling, include the **hierarchical-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level. You can optionally include the **maximum-hierarchy-levels** option. If you choose to set this option, the only supported value is 2.

```
[edit interfaces]
xe-2/0/0 {
  hierarchical-scheduler {
    maximum-hierarchy-levels 2;
  }
}
```

Three-Level Hierarchical Scheduling

Three-level hierarchical scheduling is supported only on MX Series routers running MPC/MIC interfaces. Three-level hierarchical scheduling has up to eight class of service queues. You can configure many different three-level scheduling hierarchies, depending on the location of the interface set and the use of underlying interfaces. In all variations, the physical interface is a level 1 CoS scheduler node and the queues reside at level 4.

NOTE: **Three-level hierarchical scheduling is supported only on subscriber interfaces and interface sets running over aggregated Ethernet interfaces on MPC/MIC interfaces in MX Series routers.**

When you use three-level hierarchical scheduling, interface sets can reside at either level 2 or level 3. You can also configure an underlying logic interface at level 2 and a logical interface at level 3. Table 80 on page 958 summarizes the most common cases of the interface hierarchy and the CoS scheduler node levels for three-level hierarchical scheduling.

Table 80: Three-Level Hierarchical Scheduling-Interface Hierarchy Versus CoS Scheduling Node Levels

Level 1	Level 2	Level 3	Level 4
Physical interface	Interface set	Logical interfaces	One or more queues
Physical interface	Logical interface	Interface set	One or more queues
Physical interface	Underlying logical interface	Logical interfaces	One or more queues

Copyright © 2013, Juniper Networks, Inc.

In three-level hierarchical scheduling, the CoS scheduler nodes at level 1, level 2, and level 3 form a hierarchical relationship; this differs from two-level hierarchical scheduling where no hierarchical relationship is formed.

With a three-level hierarchical scheduling, logical interfaces can reside at level 2, or they can reside at level 3, if the logical interface at level 2 is an underlying logical interface. This is shown in Figure 24 on page 959.

Figure 24: Three-Level Hierarchical Scheduling: Logical Interfaces at Level 3 with Underlying Logical Interfaces at Level 2

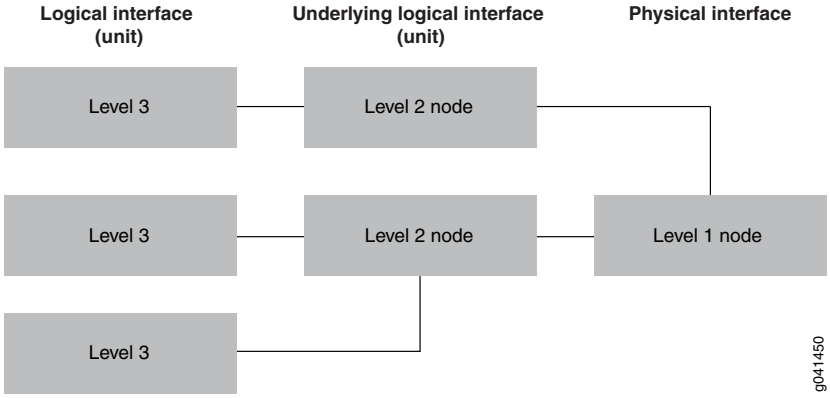

Another possible configuration for three-level hierarchical scheduling is shown in Figure 25 on page 959. In this configuration, the logical interfaces are located at level 2 and the interface sets are located at level 3.

Figure 25: Three-Level Hierarchical Scheduling: Logical Interfaces at Level 2 with Interface Sets at Level 3

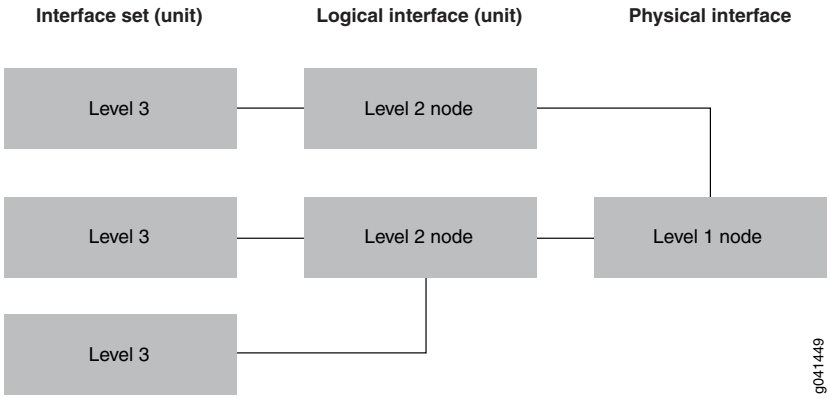

To configure three-level hierarchical scheduling, include the **implicit-hierarchy** option at the **[edit interfaces *interface-name* hierarchical-scheduler]** hierarchy level.

```
[edit interfaces]
xe-2/0/0 {
  hierarchical-scheduler {
    implicit-hierarchy;
  }
}
```

Interface Hierarchy Versus CoS Hierarchy

An interface hierarchy and a CoS scheduling hierarchy are distinctly different. Interface hierarchy refers to the relationship between the various interfaces; for example, the relationship between logical interfaces and an interface set, the relationship between a logical interface and an underlying logical interface, or the relationship between the physical interface and logical interface. CoS scheduling hierarchy refers to the hierarchical relationship between the CoS scheduler nodes. In two-level hierarchical scheduling, no hierarchy is formed between the CoS scheduler nodes; all logical interfaces and interface sets share a single level 2 scheduler node. However, when you use the **implicit-hierarchy** option for three-level hierarchical scheduling, the CoS scheduler nodes form a scheduling hierarchy.

Figure 26 on page 960 and Figure 27 on page 961 provide two scenarios for this discussion. Figure 26 on page 960 shows an interface hierarchy where a Gigabit Ethernet interface (GE-1/0/0) is the physical interface. Two logical interfaces (GE-1/0/0.100 and GE-1/0/0.101) are configured on the physical interface:

- Logical interface GE-1/0/0.100 is a member of a PPPoE interface set and a Demux interface set.
- Logical interface GE-1/0/0.101 is a member of a demux interface set.

Figure 26: Logical Interfaces at Level 2 and Interface Sets at Level 3

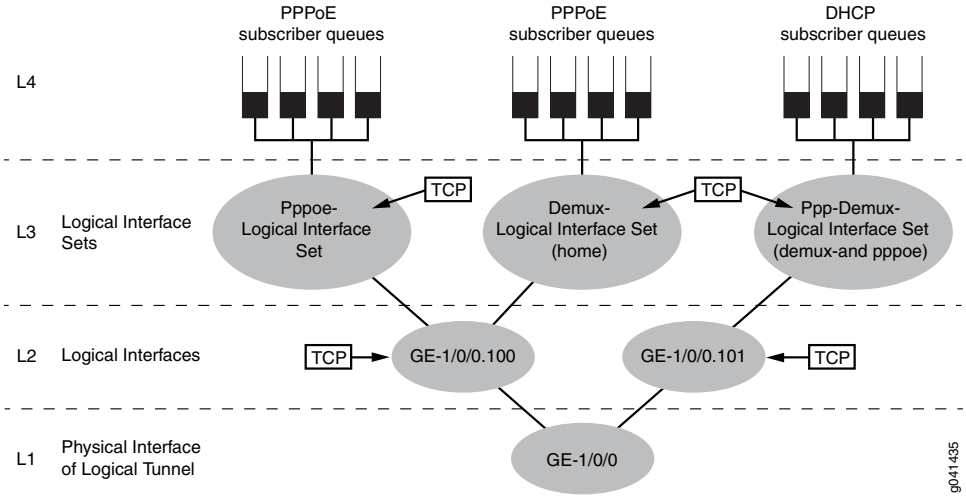

Each interface set has a dedicated queue. The CoS scheduler nodes at level 1 (physical interface), level 2 (underlying logical interfaces), and level 3 (interface sets) form a scheduling hierarchy.

To configure this scenario, you must include the **implicit-hierarchy** option under the **hierarchical-scheduler** statement on physical interface GE-1/0/0 and configure and apply traffic control profiles on each interface set and underlying logical interface.

Figure 27 on page 961 shows an interface hierarchy where Gigabit Ethernet interface GE-1/0/0 is the physical interface. Three logical interfaces are configured:

 Copyright © 2013, Juniper Networks, Inc.

- Two logical interfaces (Pp0.100 and Demux0.100) reside on the underlying logical interface GE-1/0/0.100.
- A third logical interface (Pp0.101) resides on the underlying logical interface GE-1/0/0.101.

Figure 27: Logical Interfaces at Level 3 and Underlying Logical Interfaces at Level 2

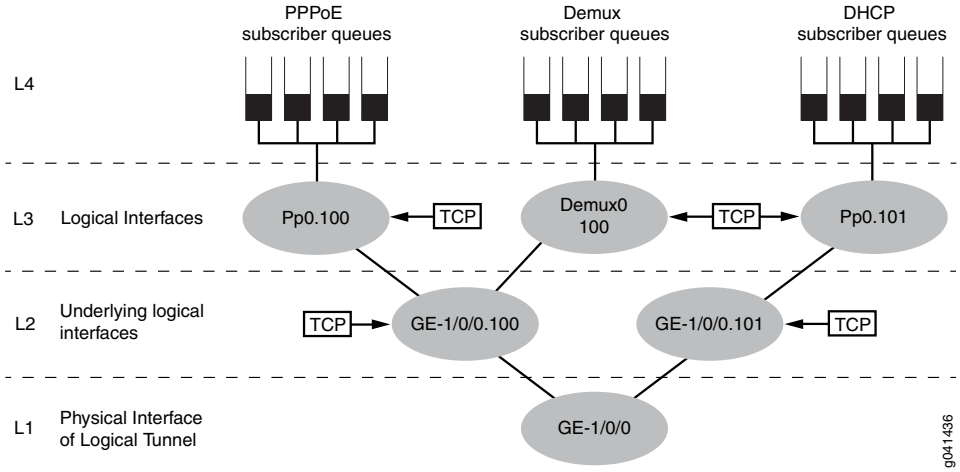

Each logical interface has a dedicated queue. The CoS scheduler nodes at level 1 (physical interface), level 2 (underlying logical interfaces), and level 3 (logical interfaces) form a scheduling hierarchy.

To configure this scenario, you must include the **implicit-hierarchy** option under the **hierarchical-scheduler** statement on physical interface GE-1/0/0 and configure and apply traffic control profiles on each logical interface and underlying logical interface.

You can configure many different three-level scheduling hierarchies; Figure 26 on page 960 and Figure 27 on page 961 present just two possible scenarios. Table 80 on page 958 summarizes the possible interface locations and CoS scheduler nodes.

Related Documentation

- *Configuring Hierarchical Schedulers for CoS*
- Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links on page 1023
- Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface on page 1024
- CoS Three-Level Hierarchical Scheduling on MPLS Pseudowire Subscriber Interfaces on page 927
- **hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) on page 1677**

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 52

Configuration Summary of Dynamic CoS for Subscriber Access

- Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 963
- Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 965
- Configuring Per-Unit Scheduling in a Dynamic Profile for Subscriber Access on page 967

Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access

You configure static scheduling and queuing in a dynamic profile for subscriber access.

To configure CoS in a dynamic profile for subscriber access using static scheduling and queuing parameters:

1. Configure the static CoS parameters in the **[edit class-of-service]** hierarchy.
   a. Enable the hierarchical scheduler for the interface.

      See "Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers" on page 956.

   b. Configure the scheduler map and schedulers.

      When you configure static scheduling and queuing in a dynamic profile, you reference the scheduler map in the dynamic profile.

      See *Configuring Schedulers*.

   c. Configure the drop profiles.

      See *Configuring RED Drop Profiles*.

   d. Configure the forwarding classes.

      See *Configuring Forwarding Classes*

   e. Configure the rewrite-rules and classifier definitions.

      See *Configuring Rewrite Rules* and *Defining Classifiers*.

Copyright © 2013, Juniper Networks, Inc.

See the *Junos OS Class of Service Library for Routing Devices* for information about configuring the remaining CoS parameters.

2. Configure a static or dynamic subscriber interface that can be referenced in the dynamic profile.
   - For static VLAN interfaces, see "Configuring Static Subscriber Interfaces in Dynamic Profiles" on page 765.
   - For dynamic VLAN interfaces, see "Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet" on page 824.
   - For dynamic IP demux interfaces, see "Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles" on page 771 and "Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet" on page 825.
   - For dynamic VLAN demux interfaces, see "Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles" on page 772.
   - For dynamic PPPoE interfaces, see "Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles" on page 899.
3. Configure CoS parameters in a dynamic profile.
   a. Configure the dynamic profile.

      See "Configuring a Basic Dynamic Profile" on page 673.

   b. Configure traffic shaping and scheduling parameters in the dynamic profile using a traffic-control profile.

      Reference the scheduler map you configured in the static **[edit class-of-service]** hierarchy.

      See "Configuring Static Traffic Shaping and Scheduling Parameters in a Dynamic Profile" on page 971.

   c. Apply CoS parameters to a subscriber interface by referencing an interface in the dynamic profile.

      See "Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile" on page 979.
4. To configure default values for subscribers on login, and enable subscribers to replace other CoS parameters when replacing services, configure variables in the dynamic profile.

   See "Configuring User-Defined CoS Variables in a Dynamic Service Profile" on page 1000.

Related Documentation

- For hardware requirements and configuration guidelines, see Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for Subscriber Access Overview on page 949
- Example: Configuring Static Hierarchical Scheduling and Queuing for Subscriber Access on page 1031

Copyright © 2013, Juniper Networks, Inc.

Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access

You can configure dynamic scheduling and queuing in dynamic profile for subscriber access.

To configure dynamic scheduling and queuing for subscriber access using dynamic scheduling and queuing parameters:

1. Configure the static CoS parameters in the **[edit class-of-service]** hierarchy.

   a. Enable the hierarchical scheduler for the interface.

      See "Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers" on page 956 and **hierarchical-scheduler**.

   b. Configure the drop profiles.

      See *Configuring RED Drop Profiles*.

   c. Configure the forwarding classes.

      See *Configuring Forwarding Classes*

   d. Configure the rewrite-rules and classifier definitions.

      See *Configuring Rewrite Rules* and *Defining Classifiers*.

   See the *Junos OS Class of Service Library for Routing Devices* for information about configuring the remaining CoS parameters.

2. Configure a static or dynamic subscriber interface that can be referenced in the dynamic profile.

   - For static VLAN interfaces, see "Configuring Static Subscriber Interfaces in Dynamic Profiles" on page 765.
   - For dynamic VLAN interfaces, see "Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet" on page 824.
   - For dynamic IP demux interfaces, see "Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles" on page 771 and "Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet" on page 825.
   - For dynamic VLAN demux interfaces, see "Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles" on page 772.
   - For dynamic PPPoE interfaces, see "Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles" on page 899.

Copyright © 2013, Juniper Networks, Inc.

3. Configure CoS parameters in a dynamic profile.

   a. Configure the dynamic profile.

      See "Configuring a Basic Dynamic Profile" on page 673.

   b. Configure traffic shaping and scheduling parameters in the dynamic profile using a traffic-control profile.

      See "Configuring Traffic Scheduling and Shaping for Subscriber Access" on page 971.

   c. Configure the schedulers and scheduler map in the dynamic profile.

      You can configure the schedulers using dynamic variables or a combination of both static values and dynamic variables.

      See "Configuring Schedulers in a Dynamic Profile for Subscriber Access" on page 973.

   d. Apply CoS parameters to a subscriber interface by referencing an interface in the dynamic profile.

      - For traffic shaping and scheduling, see "Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile" on page 979.
      - For rewrite-rules, see "Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile" on page 980.
      - For classifiers, see "Applying a Classifier to a Subscriber Interface in a Dynamic Profile" on page 981.

4. (Optional) Configure variables in access and service profiles to enable RADIUS to activate subscriber and upgrade services through CoA.

   NOTE: **Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.**

   a. Configure user-defined CoS variables in a dynamic profile.

      See "Configuring User-Defined CoS Variables in a Dynamic Service Profile" on page 1000

   b. (Optional) Enable multiple clients for the same subscriber (logical interface) to aggregate attributes by configuring the **aggregate-clients** option for the dynamic profile attached to a DHCP subscriber interface.

      See "Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces" on page 236.

      Because you have configured the scheduler map in the dynamic profile, queues are merged when subscribers change services. Other CoS parameters are replaced.

      When multiple subscribers are enabled on a DHCP subscriber interface, and the dynamic profile referenced by DHCP does not have the **replace** keyword configured, the system does not replace the parameters. Instead, it combines the values of the parameters to their maximum scalar value.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- hierarchical-scheduler on page 1677
- For hardware requirements and configuration guidelines, see Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for Subscriber Access Overview on page 949
- Example: Configuring Dynamic Hierarchical Scheduling and Queuing for Subscriber Access on page 1033

Configuring Per-Unit Scheduling in a Dynamic Profile for Subscriber Access

Per-unit scheduling enables one set of output queues for each logical interface configured under the physical interface. In per-unit scheduling configurations, each Layer 3 scheduler node is allocated a dedicated set of queues.

If you do not explicitly configure CoS parameters, a default traffic profile with queues is attached to the logical interface.

To configure per-unit scheduling and queuing for subscriber access:

1. Configure the static CoS parameters in the **[edit class-of-service]** hierarchy.

   a. Enable the per-unit scheduler for the physical interface.

      [edit interfaces *interface-name*]
      user@host# **set per-unit-scheduler**

   b. Configure the drop profiles.

      See *Configuring RED Drop Profiles*.

   c. Configure the forwarding classes.

      See *Configuring Forwarding Classes*

   d. Configure the rewrite-rules and classifier definitions.

      See *Configuring Rewrite Rules* and *Defining Classifiers*.

   See the *Junos OS Class of Service Library for Routing Devices* for information about configuring the remaining CoS parameters.

2. Configure a static or dynamic subscriber interface that can be referenced in the dynamic profile.

   - For static VLAN interfaces, see "Configuring Static Subscriber Interfaces in Dynamic Profiles" on page 765.
   - For dynamic IP demux interfaces, see "Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles" on page 771.
   - For dynamic PPPoE interfaces, see "Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles" on page 899.

3. Configure CoS parameters in a dynamic profile.

   a. Configure the dynamic profile.

      See "Configuring a Basic Dynamic Profile" on page 673.

   b. Configure traffic shaping and scheduling parameters in the dynamic profile using a traffic-control profile.

      See "Configuring Traffic Scheduling and Shaping for Subscriber Access" on page 971.

   c. Configure the schedulers and scheduler map in the dynamic profile.

      You can configure the schedulers using dynamic variables or a combination of both static values and dynamic variables.

      See "Configuring Schedulers in a Dynamic Profile for Subscriber Access" on page 973.

   d. Apply CoS parameters to a subscriber interface by referencing an interface in the dynamic profile.

      - For traffic shaping and scheduling, see "Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile" on page 979.
      - For rewrite rules, see "Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile" on page 980.
      - For classifiers, see "Applying a Classifier to a Subscriber Interface in a Dynamic Profile" on page 981.

4. (Optional) Configure variables in access and service profiles to enable RADIUS to activate subscriber and upgrade services through CoA.

   NOTE: **Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.**

   Because you have configured the scheduler map in the dynamic profile, queues are merged when subscribers change services. Other CoS parameters are replaced.

   When multiple subscribers are enabled on a DHCP subscriber interface, and the dynamic profile referenced by DHCP does not have the **replace** keyword configured, the system does not replace the parameters. Instead, it combines the values of the parameters to their maximum scalar value.

   a. Configure CoS variables in a dynamic profile.

      See "Configuring User-Defined CoS Variables in a Dynamic Service Profile" on page 1000

   b. (Optional) Enable multiple clients for the same subscriber (logical interface) to aggregate attributes by configuring the **aggregate-clients** option for the dynamic profile attached to a DHCP subscriber interface.

      See "Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces" on page 236.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- CoS for Subscriber Access Overview on page 949
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- *Example: Configuring Per-Unit Scheduling for Subscriber Access*

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 53

Configuring Dynamic Shaping and Scheduling for Subscriber Access

- Configuring Traffic Scheduling and Shaping for Subscriber Access on page 971
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973
- Applying CoS Parameters to a Subscriber Interface in a Dynamic Profile on page 979
- Configuring Scheduler and Scheduler Map Sharing on page 982
- Applying CoS Attributes to VLANs Using Agent-Circuit-Identifiers on page 984
- Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 986

Configuring Traffic Scheduling and Shaping for Subscriber Access

You use traffic-control profiles to configure traffic shaping and scheduling properties. When you reference a traffic-control profile in a dynamic profile, you can provide hierarchical shaping and scheduling for a subscriber interface.

You can choose to configure static values or dynamic variables for the shaping parameters. The values for the dynamic variables are obtained from RADIUS when a subscriber logs in or when a subscriber changes services.

You cannot configure a traffic-control profile that contains a combination of static and dynamic parameters.

This topic includes the following tasks:

- Configuring Static Traffic Shaping and Scheduling Parameters in a Dynamic Profile on page 971
- Configuring Dynamic Traffic Shaping and Scheduling Parameters in a Dynamic Profile on page 972

Configuring Static Traffic Shaping and Scheduling Parameters in a Dynamic Profile

To configure static traffic shaping and scheduling parameters in a traffic-control profile:

1. Create the traffic-control profile and assign a name.

```
[edit dynamic-profiles business-profile class-of-service]
user@host# edit traffic-control-profiles profile-name
```

2. Do one of the following:

   - Reference a static scheduler map in the dynamic profile. The scheduler map is statically configured in the **[edit class-of-service]** hierarchy.

     [edit dynamic-profiles business-profile class-of-service traffic-control-profiles *profile-name*]
     user@host# set scheduler-map *map-name*

   - Reference a dynamic scheduler map in the dynamic profile. The scheduler map is dynamically configured in the **[edit dynamic-profiles *profile-name* class-of-service scheduler-maps]** hierarchy.

     [edit dynamic-profiles business-profile class-of-service traffic-control-profiles *profile-name*]
     user@host# set scheduler-map *map-name*

3. Configure the shaping rate to be used in the dynamic profile.

   [edit dynamic-profiles business-profile class-of-service traffic-control-profiles *profile-name*]
   user@host# set shaping-rate (*rate* <burst-size *bytes*>

4. Configure the guaranteed rate to be used in the dynamic profile.

   [edit dynamic-profiles business-profile class-of-service traffic-control-profiles *profile-name*]
   user@host# set guaranteed-rate (*rate* <burst-size *bytes*>

5. Configure the delay-buffer rate.

   If you do not include this statement, the delay-buffer rate is based on the guaranteed rate if one is configured, or on the shaping rate if no guaranteed rate is configured.

   [edit dynamic-profiles business-profile class-of-service traffic-control-profiles *profile-name*]
   user@host# set delay-buffer-rate (percent *percentage* | *rate*)

Configuring Dynamic Traffic Shaping and Scheduling Parameters in a Dynamic Profile

You can configure variables for the traffic shaping and scheduling parameters. The values for the parameters are dynamically obtained by RADIUS when a subscriber logs in or changes a service.

To configure dynamic traffic-control profiles in a dynamic profile:

1. Create the traffic-control profile.

   [edit dynamic-profiles business-profile class-of-service]
   user@host# edit traffic-control-profiles *profile-name*

2. Configure the scheduler map variable.

   [edit dynamic-profiles business-profile class-of-service traffic-control-profiles *profile-name*]
   user@host# set scheduler-map $junos-cos-scheduler-map

3. Configure the shaping rate variable.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit dynamic-profiles business-profile class-of-service traffic-control-profiles
  profile-name]
user@host# set shaping-rate $junos-cos-shaping-rate <burst-size  bytes>
```

4. Configure the guaranteed rate variable.

```
[edit dynamic-profiles business-profile class-of-service traffic-control-profiles
  profile-name]
user@host# set guaranteed-rate $junos-cos-guaranteed-rate <burst-size [ bytes |
  $junos-cos-guaranteed-rate-burst]>
```

5. Configure a variable for the delay-buffer rate.

   If you do not include this statement, the delay-buffer rate is based on the guaranteed rate if one is configured, or the shaping rate if no guaranteed rate is configured.

```
[edit dynamic-profiles business-profile class-of-service traffic-control-profiles
  profile-name]
user@host# set delay-buffer-rate $junos-cos-delay-buffer-rate
```

Related Documentation

- For hardware requirements and configuration guidelines, see Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for Subscriber Access Overview on page 949
- Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 963
- Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 965
- Example: Configuring Static Hierarchical Scheduling and Queuing for Subscriber Access on page 1031
- Example: Configuring Dynamic Hierarchical Scheduling and Queuing for Subscriber Access on page 1033
- Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 986

Configuring Schedulers in a Dynamic Profile for Subscriber Access

You use schedulers to define the parameters of output queues. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the tail drop profiles associated with the queue.

You can configure up to four schedulers in a dynamic profile.

Within a dynamic profile, you can choose to define schedulers with static values, dynamic variables, or a combination of static values and dynamic variables. The dynamic variables

enable RADIUS to provide the value for the scheduler parameter when the subscriber logs in.

- Configuring Static Schedulers in a Dynamic Profile on page 974
- Configuring Dynamic Schedulers with Variables in a Dynamic Profile on page 975
- Configuring a Combination of Static and Dynamic Scheduler Parameters in a Scheduler Definition on page 976

Configuring Static Schedulers in a Dynamic Profile

This topic describes how to configure schedulers with static values in a dynamic profile for subscriber access.

To configure static scheduling and queuing in a dynamic profile:

1. Configure the scheduler and queuing parameters.

   a. Specify the scheduler for which you want to configure parameters.

      [edit dynamic-profiles *profile-name* class-of-service]
      user@host# **edit schedulers *scheduler-name***

   b. Configure the buffer size.

      [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
      user@host# **set buffer-size remainder**

   c. Configure the drop-profile map and drop profile.

      [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
      user@host# **set drop-profile-map loss-priority any protocol any drop-profile d3**

   d. Configure the priority.

      [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
      user@host# **set priority low**

   e. Configure the transmit rate.

      [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
      user@host# **set transmit-rate percent 40**

   f. Configure the excess rate.

      [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
      user@host# **set excess-rate percent 90**

   g. (Optional) Configure the priority value for the excess-rate.

      [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
      user@host# **set excess-priority high**

2. Associate the scheduler with a scheduler map.

   a. Configure the scheduler map name.

      [edit dynamic-profiles *profile-name* class-of-service]
      user@host# **set scheduler-maps data-smap**

Copyright © 2013, Juniper Networks, Inc.

b. Configure the forwarding class.

```
[edit dynamic-profiles profile-name class-of-service scheduler-maps map-name]
user@host# set forwarding-class be
```

c. Configure the scheduler.

```
[edit dynamic-profiles profile-name class-of-service scheduler-maps map-name
  forwarding-class forwarding-class-name]
user@host# set scheduler be_sch
```

Configuring Dynamic Schedulers with Variables in a Dynamic Profile

You can configure variables for the dynamic scheduler parameters. These values are dynamically obtained by RADIUS when a subscriber logs in or changes a service using a RADIUS change of authorization (CoA) message.

To configure dynamic scheduling and queuing in a dynamic profile:

1. Configure the scheduler and queuing parameters.

   a. Specify the scheduler name using a variable.

   ```
   [edit dynamic-profiles profile-name class-of-service]
   user@host# edit schedulers $junos-cos-scheduler
   ```

   b. Configure the variable for the buffer size.

   ```
   [edit dynamic-profiles profile-name class-of-service schedulers]
   user@host# set buffer-size (percent $junos-cos-scheduler-bs | temporal
     $junos-cos-scheduler-bs)
   ```

   c. Configure the variables for the drop-profile maps and the drop profile.

   ```
   [edit dynamic-profiles profile-name class-of-service schedulers]
   user@host# set drop-profile-map loss-priority low protocol any drop-profile
     $junos-cos-scheduler-low
   user@host# set drop-profile-map loss-priority medium-low protocol any
     drop-profile $junos-cos-scheduler-medium-low
   user@host# set drop-profile-map loss-priority medium-high protocol any
     drop-profile $junos-cos-scheduler-medium-high
   user@host# set drop-profile-map loss-priority high protocol any drop-profile
     $junos-cos-scheduler-high
   user@host# set drop-profile-map loss-priority any protocol any drop-profile
     $junos-cos-scheduler-any
   ```

   d. Configure the variable for the priority.

   ```
   [edit dynamic-profiles profile-name class-of-service schedulers]
   user@host# set priority $junos-cos-scheduler-pri
   ```

   e. Configure the variable for the transmit rate.

   ```
   [edit dynamic-profiles profile-name class-of-service schedulers]
   user@host# set transmit-rate $junos-cos-scheduler-tx
   ```

   f. Configure the variable for the excess rate.

   ```
   [edit dynamic-profiles profile-name class-of-service schedulers]
   ```

```
user@host# set excess-rate percent $junos-cos-scheduler-excess-rate
```

g. Configure the variable for the priority of the excess-rate.

```
[edit dynamic-profiles profile-name class-of-service schedulers]
user@host# set excess-priority $junos-cos-scheduler-excess-priority
```

2. Associate the scheduler with a scheduler map.

a. Configure the scheduler map name.

```
[edit dynamic-profiles profile-name class-of-service]
user@host# edit scheduler-maps scheduler-map-name
```

b. Configure the forwarding class.

```
[edit dynamic-profiles profile-name class-of-service scheduler-maps
  scheduler-map-name]
user@host# set forwarding-class be
```

c. Configure the scheduler.

```
[edit dynamic-profiles profile-name class-of-service scheduler-maps
  scheduler-map-name]
user@host# set scheduler $junos-cos-scheduler
```

Configuring a Combination of Static and Dynamic Scheduler Parameters in a Scheduler Definition

Within a dynamic profile, you can choose to configure one dynamic scheduler definition, or combine static and dynamic scheduler parameters in many static scheduler definitions.

Combining static and dynamic scheduler parameters enables you to provide subscribers with unique rate configurations that the RADIUS definitions for predefined variables do not allow.

To configure a scheduler definition that contains static and dynamic scheduling and queuing parameters:

1. Configure the scheduler definition.

a. Specify the scheduler name.

NOTE: To configure a static scheduler that contains both static and dynamic parameters, you must specify a unique scheduler name, not the $junos-cos-scheduler variable.

```
[edit dynamic-profiles profile-name class-of-service]
user@host# edit schedulers scheduler-name
```

b. Configure the buffer size.

Do either of the following:

- Configure a static value.

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
```


Copyright © 2013, Juniper Networks, Inc.

```
user@host# set buffer-size (percent percentage | remainder | temporal
  (microseconds)
```

- Configure a variable.

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
user@host# set buffer-size (percent $junos-cos-scheduler-bs | temporal
  $junos-cos-scheduler-bs)
```

c. Configure the drop-profile maps, the drop profile, and the priority.

Do either of the following:

- Configure static values.

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
user@host# set drop-profile-map loss-priority any protocol any drop-profile d3
```

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
user@host# set priority low
```

- Configure variables.

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
user@host# set drop-profile-map loss-priority low protocol any drop-profile
  $junos-cos-scheduler-low
user@host# set drop-profile-map loss-priority medium-low protocol any
  drop-profile $junos-cos-scheduler-medium-low
user@host# set drop-profile-map loss-priority medium-high protocol any
  drop-profile $junos-cos-scheduler-medium-high
user@host# set drop-profile-map loss-priority high protocol any drop-profile
  $junos-cos-scheduler-high
user@host# set drop-profile-map loss-priority any protocol any drop-profile
  $junos-cos-scheduler-any
```

d. Configure the priority.

Do either of the following:

- Configure a static value.

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
user@host# set excess-priority high
```

- Configure a variable.

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
user@host# set excess-priority $junos-cos-scheduler-excess-priority
```

e. Configure the transmit rate.

Do either of the following:

- Configure a static value.

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
user@host# set transmit-rate
```

- Configure a variable.

```
[edit dynamic-profiles profile-name class-of-service schedulers scheduler-name]
user@host# set transmit-rate $junos-cos-scheduler-tx
```

f. Configure the excess rate.

Do either of the following:

- Configure a static value.

  [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
  user@host# **set** excess-rate **percent 250**

- Configure a variable.

  [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
  user@host# **set** excess-rate **percent $junos-cos-scheduler-excess-rate**

g. Configure the priority for the excess-rate.

Do either of the following:

- Configure a static value.

  [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
  user@host# **set** excess-priority **high**

- Configure a variable.

  [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
  user@host# **set** excess-priority **percent $junos-cos-scheduler-excess-priority**

2. Associate the scheduler with a scheduler map.

a. Configure the scheduler map name.

[edit dynamic-profiles *profile-name* class-of-service]
user@host# **edit** scheduler-maps ***scheduler-map-name***

b. Configure the forwarding class.

[edit dynamic-profiles *profile-name* class-of-service scheduler-maps *scheduler-map-name*]
user@host# **set** forwarding-class **be**

c. Configure the scheduler.

[edit dynamic-profiles *profile-name* class-of-service scheduler-maps *scheduler-map-name*]
user@host# **set** scheduler **$junos-cos-scheduler**

Related Documentation

- For hardware requirements and configuration guidelines, see Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 965
- Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 986
- Changing CoS Services Overview on page 993

Copyright © 2013, Juniper Networks, Inc.

Applying CoS Parameters to a Subscriber Interface in a Dynamic Profile

You provide CoS parameters to a subscriber by associating the CoS parameters with an interface in a dynamic profile.

Traffic and scheduling parameters can be configured in the dynamic profile and associated with a subscriber by attaching an output traffic control profile to the interface in the dynamic profile.

You configure rewrite rules and classifiers statically in the **[edit class-of-service]** hierarchy and reference them in the dynamic profile.

- Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile on page 979
- Applying Minimal Shaping and Scheduling to Remaining Subscriber Traffic on page 980
- Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile on page 980
- Applying a Classifier to a Subscriber Interface in a Dynamic Profile on page 981

Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile

After you configure the traffic shaping and scheduling CoS parameters in a dynamic profile, you apply them to an interface. The output-traffic control profile enables you to provide traffic scheduling to the interface.

To apply CoS attributes to an interface in a dynamic profile:

1. Specify that you want to apply CoS attributes to an interface in the dynamic profile.

   user@host# **edit** dynamic-profiles ***profile-name*** class-of-service

2. Configure the interface name and logical interface using a variable, and apply the output-traffic control profile to the interface.

   [edit dynamic-profiles *profile-name* class-of-service interfaces]
   user@host# **set** interfaces **$junos-interface-ifd-name** unit **$junos-underlying-interface-unit** output-traffic-control-profile ***profile-name***

   You can use one of the following methods to specify the output-traffic control profile you want to use:

   - Reference the **$junos-cos-traffic-control-profile** predefined variable. At subscriber login, subscriber management takes one of the following actions, in the order listed:

     a. If RADIUS is being used and it returns a value for the traffic-control profile, subscriber management uses the RADIUS value.

     b. If RADIUS is not being used, subscriber management uses the default traffic-control profile (which is specified by the **predefined-variables-default** statement at the **[edit dynamic-profiles]** hierarchy).

        For example:

```
user@host# set interfaces $junos-interface-ifd-name unit
  $junos-underlying-interface-unit output-traffic-control-profile
  $junos-cos-traffic-control-profile
```

- Explicitly reference the name of the traffic-control profile.

  For example:

```
user@host# set interfaces $junos-interface-ifd-name unit
  $junos-underlying-interface-unit output-traffic-control-profile tcp-sales-2
```

Applying Minimal Shaping and Scheduling to Remaining Subscriber Traffic

It is beneficial to apply a remaining traffic-control profile to a logical interface to provide minimal CoS scheduling when you have not configured or over-provisioned Layer 3 schedulers. In the event that schedulers are not available, the remaining subscriber traffic receives the essential level of service.

To configure scheduling for remaining subscriber traffic:

1. Enable hierarchical scheduling for the interface.

   [edit interfaces *interface-name*]
   user@host# set hierarchical-scheduler

2. Apply the remaining traffic control profile to the port on which you enabled hierarchical scheduling.

   [edit class-of-service interfaces *interface-name*]
   user@host# set output-traffic-control-profile-remaining *profile-name*

Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile

Rewrite rules define the marking for various CoS values, including DSCP, DSCP IPv6, IP precedence, and IEEE 802.1 CoS values. Rewrite rules have an associated forwarding class and code-point alias or bit set.

NOTE: By default, subscriber lawful intercept does not intercept DHCP control packets that are generated by the routing engine. To ensure that a DHCP control packet generated by the routing engine is intercepted, you need to configure the ieee-802.1 rewrite-rule for VLAN demux.

For dynamic CoS, you define the rewrite rules mapping for the CoS values statically, then reference the rewrite rule configuration in the dynamic profile for the subscriber interface.

To configure a rewrite rule in a dynamic profile:

1. Define the rewrite-rules mapping for the traffic that passes through all queues on the interface. The available rewrite-rules types for dynamic CoS are **dscp**, **dscpv6**, **ieee-802.1** and **inet-precedence**.

   See *Configuring Rewrite Rules*.

2. Apply the rewrite-rules definition to the subscriber interface in the dynamic profile.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
  logical-unit-number]
user@host# edit rewrite-rules
```

3. Configure the applicable rewrite rule markers in the dynamic profile.

   - For DSCP:

     ```
     [edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
       logical-unit-number rewrite-rules]
     user@host# set dscp (rewrite-name | default)
     ```

   - For DSCPv6:

     ```
     [edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
       logical-unit-number rewrite-rules]
     user@host# set dscp-ipv6 (rewrite-name | default)
     ```

   - For IEEE 802.1:

     ```
     [edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
       logical-unit-number rewrite-rules]
     user@host# set ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner)
     ```

   - For inet-precedence:

     ```
     [edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
       logical-unit-number rewrite-rules]
     user@host# set inet-precedence (rewrite-name | default)
     ```

Applying a Classifier to a Subscriber Interface in a Dynamic Profile

You can apply the classification map to a subscriber interface in a dynamic profile.

For dynamic CoS, you define the classification map for the CoS values statically, then reference the classifier configuration in the dynamic profile for the subscriber interface.

To apply a classifier to an interface in a dynamic profile:

1. Define the classifier.

   The available classifier types for dynamic CoS are **dscp**, **dscp-ipv6**, **ieee-802.1**, and **inet-precedence**.

   See *Defining Classifiers*.

2. Apply the classifier definition to the subscriber interface in the dynamic profile.

   ```
   [edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
     logical-unit-number]
   user@host# edit classifiers
   ```

3. Configure the applicable classifiers in the dynamic profile.

   - For DSCP:

     ```
     [edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
       logical-unit-number classifiers]
     user@host# set dscp (classifier-name | default)
     ```

   - For DSCPv6:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
  logical-unit-number classifiers]
user@host# set dscp-ipv6 (classifier-name | default)
```

- For IEEE 802.1:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
  logical-unit-number classifiers]
user@host# set ieee-802.1 (classifier-name | default) vlan-tag (inner | outer)
```

- For inet-precedence:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name unit
  logical-unit-number classifiers]
user@host# set inet-precedence (classifier-name | default)
```

Related Documentation

- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 963
- Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 965
- CoS for Subscriber Access Overview on page 949

Configuring Scheduler and Scheduler Map Sharing

The system generates unique identifiers (IDs) in dynamic profiles created for services. The generated unique IDs enable you to identify and configure separate parameter values with the same variable name. When applied to CoS, you can now configure scheduler and scheduler map sharing. In client-access profiles, schedulers and scheduler maps must use the unique ID format. If the client-access profile uses the unique ID format and you want to have either scheduler or scheduler map sharing for service activation, you must configure the service profile in unique ID format. Generating unique IDs based on schedulers and scheduler maps eliminates duplication and improves router performance and scalability. You can configure scheduler and scheduler map sharing by including the variables for CoS in the client access or service dynamic profile. All scheduler maps and schedulers must be in the unique ID format.

Before you configure variables for the client access or service dynamic profile:

- Create a basic dynamic profile.

  See "Configuring a Basic Dynamic Profile" on page 673.

To configure variables for the client access or service dynamic profile:

1. Configure the variables for the dynamic client access profile.

```
[edit dynamic-profiles client-profile variables]
user@host# set smap_data uid
user@host# set data_sched uid
```

2. Configure the CoS parameters for the variables in the scheduler profile.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit dynamic-profiles client-profile class-of-service]
user@host# edit schedulers “$data_sched”
user@host# set transmit-rate percent 10
user@host# set buffer-size remainder
user@host# set priority low
```

3. Configure the CoS parameters for the variables in the scheduler maps profile.

```
[edit dynamic-profiles client-profile class-of-service]
user@host# edit scheduler-maps “$smap_data”
user@host# edit forwarding-class be scheduler “$data_sched”
```

For example, you can configure scheduler maps and schedulers for a client access profile:

```
dynamic-profiles {
  cos-para {
    variables {
      data_smap uid;
      data_video_smap uid;
      voice_smap uid;
      data_sched uid;
      video_sched uid;
      voice_sched uid;
    }
    ...
    class-of-service {
      traffic-control-profiles {
        tcp1 {
          scheduler-map "$junos-cos-scheduler-map";
          shaping-rate "$junos-cos-shaping-rate";
          guaranteed-rate 10m;
          delay-buffer-rate "$junos-cos-delay-buffer-rate";
        }
      }
      interfaces {
        "$junos-interface-ifd-name" {
          unit "$junos-underlying-interface-unit" {
            output-traffic-control-profile tcp1;
          }
        }
      }
      scheduler-maps {
        "$data_smap" {
          forwarding-class be scheduler "$data_sched";
        }
        "$data_video_smap" {
          forwarding-class be scheduler "$data_sched";
          forwarding-class af scheduler "$video_sched";
        }
        “$voice_smap” {
          forwarding-class ef scheduler “$voice_sched”;
        }
      }
      schedulers {
        "$data_sched" {
          transmit-rate "$junos-cos-scheduler-tx";
          inactive: buffer-size percent "$junos-cos-scheduler-bs";
```

```
            priority "$junos-cos-scheduler-pri";
          }
          "$video_sched" {
            transmit-rate "$junos-cos-scheduler-tx";
            inactive: buffer-size percent "$junos-cos-scheduler-bs";
            priority "$junos-cos-scheduler-pri";
          }
          “$voice_sched” {
            transmit-rate percent 10;
            buffer-size remainder;;
            priority low;
          }
        }
      }
    }
  }
```

Related Documentation

- Access Profiles and Service Profiles Overview on page 670
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950

Applying CoS Attributes to VLANs Using Agent-Circuit-Identifiers

To apply CoS attributes, such as shaping, at the household level, you must set and define the CoS policy for the agent-circuit-identifier VLAN interface set using the dynamic profile for the agent-circuit-identifier interface set (not the subscriber profile). You can also configure a traffic control profile and a remaining traffic control profile for a dynamic interface set.

The following example is a CoS profile for an ACI set using a unique-ID based dynamic scheduler map:

Before you apply CoS attributes to VLANs:

- Create a basic dynamic profile.

  See “Configuring a Basic Dynamic Profile” on page 673.

Configure a CoS dynamic profile with a simple traffic control profile that is applied to the dynamic interface set that represents the ACI VLAN.

1. Configure CoS to support a dynamic interface set in the CoS profile:

   [edit dynamic-profiles *profile-name*]
   user@host# edit interface "$junos-interface-name"

2. Configure the interfaces.

   [edit dynamic-profiles *profile-name* interfaces]
   user@host# edit interface-set "$junos-interface-set-name"
   user@host# edit interface "$junos-interface-ifd-name"

3. Configure the CoS traffic control profile.

   [edit class-of-service]
   user@host# edit traffic-control-profiles *traffic-control-profile-name*

 Copyright © 2013, Juniper Networks, Inc.

```
user@host# set shaping-rate rate
user@host# set guaranteed-rate rate
```

4. Specify the interfaces.

```
[edit class-of-service interfaces]
user@host# edit interface-set "$junos-interface-set-name"
user@host# edit output-traffic-control-profile profile-name
```

The following example is a CoS profile for an ACI set using a unique ID-based dynamic scheduler map:

```
aci-set-profile {
  variables {
    ds1q0q2DP uid;
    ds1q1q2DP uid;
    be1_dp uid;
    ef1_dp uid;
    af1_dp uid;
    nc1_dp uid;
  }
  interfaces {
    interface-set "$junos-interface-set-name" {
      interface "$junos-interface-ifd-name";
    }
  }
  class-of-service {
    traffic-control-profiles {
      tcp2 {
        inactive: scheduler-map ss1q0q1DP;
        shaping-rate 50m;
        guaranteed-rate 30m;
        overhead-accounting bytes -20;
      }
      tcp3 {
        scheduler-map "$ds1q1q2DP";
        shaping-rate 30m;
        guaranteed-rate 10m;
        overhead-accounting bytes -20;
      }
    }
    interfaces {
      interface-set "$junos-interface-set-name" {
        output-traffic-control-profile tcp2;
        output-traffic-control-profile-remaining tcp3;
      }
    }
    scheduler-maps {
      "$ds1q0q2DP" {
        forwarding-class be scheduler "$be1_dp";
        forwarding-class af scheduler "$af1_dp";
        forwarding-class nc scheduler "$nc1_dp";
      }
      "$ds1q1q2DP" {
        forwarding-class ef scheduler "$ef1_dp";
        forwarding-class af scheduler "$af1_dp";
```

```
            forwarding-class nc scheduler "$nc1_dp";
          }
        }
        schedulers {
          "$be1_dp" {
            transmit-rate percent 25;
            priority low;
            drop-profile-map loss-priority low protocol any drop-profile d3;
            drop-profile-map loss-priority medium-low protocol any drop-profile d2;
            drop-profile-map loss-priority medium-high protocol any drop-profile d1;
            drop-profile-map loss-priority high protocol any drop-profile d0;
          }
          "$ef1_dp" {
            transmit-rate percent 25;
            priority low;
            drop-profile-map loss-priority low protocol any drop-profile d3;
            drop-profile-map loss-priority medium-low protocol any drop-profile d2;
            drop-profile-map loss-priority medium-high protocol any drop-profile d1;
            drop-profile-map loss-priority high protocol any drop-profile d0;
          }
          "$af1_dp" {
            transmit-rate percent 25;
            priority low;
            drop-profile-map loss-priority low protocol any drop-profile d3;
            drop-profile-map loss-priority medium-low protocol any drop-profile d2;
            drop-profile-map loss-priority medium-high protocol any drop-profile d1;
            drop-profile-map loss-priority high protocol any drop-profile d0;
          }
          "$nc1_dp" {
            transmit-rate percent 25;
            priority low;
            drop-profile-map loss-priority low protocol any drop-profile d3;
            drop-profile-map loss-priority medium-low protocol any drop-profile d2;
            drop-profile-map loss-priority medium-high protocol any drop-profile d1;
            drop-profile-map loss-priority high protocol any drop-profile d0;
          }
        }
      }
    }
```

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Changing CoS Services Overview on page 993

Verifying the Scheduling and Shaping Configuration for Subscriber Access

Purpose View the class-of-service (CoS) configurations that are referenced in a dynamic profile for subscriber access.

Action

- To display the entire CoS configuration, including static and dynamic parameters:

  user@host> **show class-of-service**

- To display the CoS configuration for a subscriber interface:

  user@host> **show class-of-service interface**

 Copyright © 2013, Juniper Networks, Inc.

- To display traffic shaping and scheduling profiles:

  user@host> **show class-of-service traffic-control-profile**

- To display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry:

  user@host> **show class-of-service scheduler-map**

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 54

RADIUS and Dynamic CoS Overview

- Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS on page 989
- Changing CoS Services Overview on page 993
- CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview on page 996
- Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions on page 998

Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS

You can configure interface-specific CoS parameters that the router obtains when subscribers log in at appropriately configured static or dynamic subscriber interfaces. This feature is supported only for interfaces on Enhanced Queuing Dense Port Concentrators (EQ DPCs) in MX Series 3D Universal Edge Routers.

To configure a dynamic profile to provide initial CoS Services, make sure you understand the following concepts:

- Dynamic Configuration of Initial CoS in Access Profiles on page 989
- Predefined Variables for Dynamic Configuration of Initial Traffic Shaping on page 990
- Predefined Variables for Dynamic Configuration of Initial Scheduling and Queuing on page 991

Dynamic Configuration of Initial CoS in Access Profiles

When a router interface receives a join message from a DHCP subscriber, the Junos OS applies the values configured in the dynamic profile associated with that router interface. A dynamic profile that is activated through its association with a subscriber interface is known as an *access dynamic profile*. You can associate a dynamic profile with a subscriber interface on the router by including statements at the **[edit dynamic-profiles *profile-name* class-of-service interfaces]** hierarchy level.

The Junos OS supports predefined variables for obtaining a scheduler-map name and traffic-shaping parameters from the RADIUS authentication server and predefined variables for obtaining a scheduler name and scheduler parameters from the RADIUS authentication server. When a client authenticates over a router interface associated

with the access dynamic profile, the router replaces the predefined variables with interface-specific values obtained from the RADIUS server.

NOTE: To associate dynamically configured initial CoS features with a subscriber interface, reference *Junos OS predefined variables*—and not *user-defined variables*—in an *access* dynamic profile for that interface.

Predefined Variables for Dynamic Configuration of Initial Traffic Shaping

You can configure an access dynamic profile that provides initial traffic-shaping parameters when a subscriber logs in. The Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.

If you define the Juniper Networks authentication and authorization VSA for CoS traffic-shaping parameter values (attribute number 26–108) on the RADIUS authentication server, the RADIUS server includes the values in RADIUS Access-Accept messages it sends to the router when a subscriber successfully authenticates over the interface.

To provide an initial scheduler map name and traffic shaping parameters obtained from the RADIUS authentication server when a subscriber logs in, reference the Junos OS predefined variables for CoS listed in Table 81 on page 990 in an access dynamic profile associated with the subscriber interface.

Table 81: CoS Predefined Variables for Scheduler Map and Traffic Shaping

Variable	Description
$junos-cos-scheduler-map	Scheduler-map name to be dynamically configured in a traffic-control profile in the access dynamic profile when a subscriber logs in. **NOTE:** The scheduler map referenced by the scheduler-map statement can be defined dynamically (at the **[edit dynamic-profiles *profile-name* class-of-service scheduler-maps]** hierarchy level) or statically (at the **[edit class-of-service scheduler-maps]** hierarchy level).
$junos-cos-shaping-rate	Shaping rate to be dynamically configured in a traffic-control profile in the access dynamic profile when a subscriber logs in. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.
$junos-cos-guaranteed-rate	Guaranteed rate to be dynamically configured in a traffic-control profile in the access dynamic profile when a subscriber logs in. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.
$junos-cos-delay-buffer-rate	Delay-buffer rate to be dynamically configured in a traffic-control profile in the access dynamic profile when a subscriber logs in. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.

Copyright © 2013, Juniper Networks, Inc.

Predefined Variables for Dynamic Configuration of Initial Scheduling and Queuing

You can configure an access dynamic profile that provides initial traffic-shaping parameters when a subscriber logs in. The Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.

If you define the Juniper Networks authentication and authorization VSA for CoS scheduling and queuing parameter values (attribute number 26–146) on the RADIUS authentication server, the RADIUS server includes the values in RADIUS Access-Accept messages it sends to the router when a subscriber successfully authenticates over the interface.

To provide an initial scheduler name and scheduler and queuing parameters obtained from the RADIUS authentication server when a subscriber logs in, reference the Junos OS predefined variables listed in Table 82 on page 991 in an access dynamic profile associated with the subscriber interface.

Table 82: CoS Predefined Variables for Scheduling and Queuing

Variable	Description
$junos-cos-scheduler	Name of a scheduler to be dynamically configured in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.
$junos-cos-scheduler-transmit-rate	Transmit rate to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.
$junos-cos-scheduler-bs	Buffer size, as a percentage of total buffer, to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.
$junos-cos-scheduler-pri	Packet-scheduling priority value to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached.
$junos-cos-scheduler-dropfile-low	Name of the drop profile for RED for loss-priority level **low** to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached. **NOTE:** The drop profile must be configured statically (at the **[edit class-of-service schedulers *scheduler-name* drop-profiles]** hierarchy level) for loss-priority **low**.

Table 82: CoS Predefined Variables for Scheduling and Queuing *(continued)*

Variable	Description
$junos-cos-scheduler-dropfile-medium-low	Name of the drop profile for RED for loss-priority level **medium-low** to be dynamically configured for the scheduler in the access dynamic profile. The Junos OS obtains this information from the RADIUS server when a subscriber authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached. **NOTE:** The drop profile must be configured statically (at the **[edit** class-of-service schedulers ***scheduler-name*** **drop-profiles]** hierarchy level).
$junos-cos-scheduler-dropfile-medium-high	Name of the drop profile for RED for loss-priority level **medium-high** to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached. **NOTE:** The drop profile must be configured statically (at the **[edit** class-of-service schedulers ***scheduler-name*** **drop-profiles]** hierarchy level).
$junos-cos-scheduler-dropfile-high	Name of the drop profile for RED for loss-priority level **high** to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached. **NOTE:** The drop profile must be configured statically (at the **[edit** class-of-service schedulers ***scheduler-name*** **drop-profiles]** hierarchy level).
$junos-cos-scheduler-dropfile-any	Name of the drop profile for RED for loss-priority level **any** to be dynamically configured for the scheduler in the access dynamic profile. You can configure a RADIUS authentication server to include this information in the Accept-Accept message when a subscriber successfully authenticates over the static or dynamic subscriber interface to which the access dynamic profile is attached. **NOTE:** The drop profile must be configured statically (at the **[edit** class-of-service schedulers ***scheduler-name*** **drop-profiles]** hierarchy level).

Related Documentation

- Subscriber Activation and Service Management in an Access Network on page 9
- Dynamic Profiles Overview on page 638
- Dynamic Variables Overview on page 640
- Junos OS Predefined Variables on page 642
- Configuring Initial CoS Parameters Dynamically Obtained from RADIUS on page 999
- Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS on page 1040

Copyright © 2013, Juniper Networks, Inc.

Changing CoS Services Overview

This topic describes how to provide CoS when subscribers dynamically upgrade or downgrade services in an access environment.

You can configure your network with an *access profile* that provides all subscribers with default CoS parameters when they log in. For example, all subscribers can receive a basic data service. By configuring the access profile with Junos OS predefined variables for RADIUS-provided CoS parameters, you also enable the service to be activated for those subscribers at login.

To enable subscribers to activate a service or upgrade to different services through RADIUS change-of-authorization (CoA) messages after login, configure a *service profile* that includes user-defined variables.

Types of CoS Variables Used in a Service Profile

You can configure variables for the following CoS parameters in a service profile:

- Shaping rate
- Delay buffer rate
- Guaranteed rate
- Scheduler map

For each CoS parameter, you must associate a RADIUS vendor ID. For each vendor ID, you must assign an attribute number and a tag. The tag is used to differentiate between values for different CoS variables when you specify the same attribute number for those variables. These values are matched with the values supplied by RADIUS during subscriber authentication. All of the values in the dynamic profile must be defined in RADIUS or none of the values are passed.

Optionally, you can configure default values for each parameter. Configuring default values is beneficial if you do not configure RADIUS to enable service changes. During service changes, RADIUS takes precedence over the default value that is configured.

Static and Dynamic CoS Configurations

Depending on how you configure CoS parameters in the access and service profiles, certain CoS parameters are replaced or merged when subscribers change or activate new services.

Static configuration is when you configure the scheduler map and schedulers in the static **[edit class-of-service]** hierarchy and reference the scheduler map in the dynamic profile. Dynamic configuration is when you configure the scheduler map and schedulers within the dynamic profile.

The CoS configuration also depends on whether you have enabled multiple subscribers on the same logical interface using the **aggregate-clients** statements in the dynamic profile referenced by DHCP. When you specify the **aggregate-clients replace** statement, the scheduler map names are replaced. In both cases, if the length of the scheduler map

Copyright © 2013, Juniper Networks, Inc.

name exceeds 128 characters, subscribers cannot log in. When you specify the **aggregate-clients merge** statement, the scheduler map names specified in the dynamic profile are appended.

BEST PRACTICE: To improve CoS performance in IPv4, IPv6, and dual-stack networks, we recommend that you use the **aggregate-clients replace** statement rather than the **aggregate-clients merge** statement.

Scenarios for Static and Dynamic Configuration of CoS Parameters

Table 83 on page 994 lists the scenarios for static and dynamic configuration of CoS parameters in access profiles and service profiles at subscriber login. The table also lists the behavior for each configuration for service activation and service modification using RADIUS CoA messages.

Table 83: CoS Services and Variables

Scenario	Static CoS Configuration (Single Subscriber)	Dynamic CoS Configuration (Single Subscriber)	Dynamic CoS Configuration (Multiple Subscribers Enabled on a Logical Interface with the **aggregate-clients merge** Statement)	Dynamic CoS Configuration (Multiple Subscribers Enabled on a Logical Interface with the **aggregate-clients replace** Statement)
Subscriber login	• Configure RADIUS values or default values for all parameters in access profile • Configure scheduler map in **edit class-of-service** hierarchy and reference in access profile	• Configure RADIUS values or default values for all parameters in access profile • Configure scheduler map and schedulers in access profile	• Configure RADIUS values or default values for all parameters in access profile • Configure scheduler map and schedulers in access profile	• Configure RADIUS values or default values for all parameters in access profile • Configure scheduler map and schedulers in access profile
RADIUS CoA for service or variable change	Replaces the following parameters: • Delay buffer rate • Guaranteed rate • Scheduler map • Shaping rate	Replaces the following parameters: • Delay buffer rate • Guaranteed rate • Shaping rate • Scheduler map	Combines the values of the following parameters to their maximum scalar value: • Delay buffer rate • Guaranteed rate • Shaping rate Appends the scheduler map parameter	Replaces the following parameters: • Delay buffer rate • Guaranteed rate • Shaping rate • Scheduler map

Copyright © 2013, Juniper Networks, Inc.

Table 83: CoS Services and Variables *(continued)*

Scenario	Static CoS Configuration (Single Subscriber)	Dynamic CoS Configuration (Single Subscriber)	Dynamic CoS Configuration (Multiple Subscribers Enabled on a Logical Interface with the aggregate-clients merge Statement)	Dynamic CoS Configuration (Multiple Subscribers Enabled on a Logical Interface with the aggregate-clients replace Statement)
RADIUS CoA for service activation	Does not merge queues **NOTE:**In this case, use a similar configuration to the access profile, including the same name for the traffic-control-profile. During service activation, this configuration replaces the original configuration in the access profile.	Merge queues if the queue specified in the service profile is not already in use for the subscriber **NOTE:** Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.	Merge queues if the queue specified in the service profile is not already in use for the subscriber **NOTE:** Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.	Merge queues if the queue specified in the service profile is not already in use for the subscriber **NOTE:** Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.

Related Documentation

- Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 963
- Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 965
- Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview on page 207
- RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework on page 90
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950

Copyright © 2013, Juniper Networks, Inc.

CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview

To control bandwidth at a household level in a subscriber access network, you can apply RADIUS dynamic class of service (CoS) traffic-shaping attributes to a dynamic interface set and its member subscriber sessions when the subscriber sessions are authenticated. (The dynamic interface set itself does not go through the authentication process.)

A *household* is represented by either a dynamic interface set or a dynamic agent-circuit-identifier (ACI) interface set from which the subscriber sessions originate. For this feature, dynamic interface sets and dynamic ACI interface sets are mapped to Level 2 of the Junos OS CoS scheduler hierarchy, which enables you to use CoS traffic-shaping to shape the bandwidth at the household (interface set) level.

The *subscriber sessions*, also referred to as *subscriber interfaces* or *client sessions*, can be dynamic VLAN, PPPoE, or IP demultiplexing (IP demux) subscriber interfaces. The subscriber interfaces are mapped to Level 3 of the Junos OS CoS scheduler hierarchy.

- Supported Network Configurations on page 996
- Traffic-Control Profiles in Subscriber Interface Dynamic Profiles on page 996
- CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets and Member Subscriber Sessions on page 997

Supported Network Configurations

Applying RADIUS dynamic CoS traffic-shaping attributes to a dynamic interface set and its member subscriber sessions is supported for the following network configurations:

- Dynamic IP demux subscriber interfaces (for DHCP subscribers) over either a dynamic interface set or a dynamic ACI interface set
- Dynamic PPPoE subscriber interfaces over either a dynamic interface set or a dynamic ACI interface set

Traffic-Control Profiles in Subscriber Interface Dynamic Profiles

To apply dynamic CoS traffic-shaping attributes to a dynamic interface set and its member subscriber sessions, you must define and attach the traffic-control profiles for *both* the dynamic interface set and the dynamic subscriber sessions within the dynamic profile for the subscriber interface.

At the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles]** hierarchy level in the dynamic profile, configure both of the following:

- Traffic-control profile for the dynamic VLAN, PPPoE, or IP demux subscriber interfaces
- Traffic-control profile for the dynamic interface set or dynamic ACI interface set to which the subscriber interfaces belong

Copyright © 2013, Juniper Networks, Inc.

RADIUS tag values for the Junos OS CoS traffic shaping predefined variables used in both traffic-control profiles must be in the 100s range, as described in "CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets" on page 1005.

At the **[edit dynamic-profiles *profile-name* interfaces]** hierarchy level in the dynamic profile, use the **output-traffic-control-profile** statement to apply the traffic-control profiles to the dynamic subscriber interface and the dynamic interface set or dynamic ACI interface set.

CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets and Member Subscriber Sessions

The set of **$junos-cos-*parameter*** predefined dynamic variables has been duplicated and assigned a RADIUS tag value in the 100s range for use with this feature. The RADIUS tag value is the only difference between the existing CoS traffic-shaping predefined dynamic variables and the predefined dynamic variables that you must use with this feature.

Both RADIUS instances of the **$junos-cos-*parameter*** predefined dynamic variables are available, but you must use the dynamic variables with tag values in the 100s range to apply CoS traffic-shaping attributes to both the dynamic interface set and member subscriber sessions in a subscriber interface dynamic profile.

For example, the existing **$junos-cos-shaping-rate** predefined variable is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 2. To apply CoS traffic-shaping attributes to the dynamic interface set and its member subscriber sessions, you must instead use the **$junos-cos-shaping-rate** predefined variable that is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 102.

NOTE: Do not configure a combination of $junos-cos-*parameter* predefined dynamic variables with RADIUS tag values in the 100s range and $junos-cos-*parameter* predefined dynamic variables with tag values not in the 100s range in the same traffic-control profile. If you do so, the subscriber authentication process fails.

Related Documentation

- Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions on page 998
- Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions on page 1003
- CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets on page 1005

Copyright © 2013, Juniper Networks, Inc.

Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions

Observe the following guidelines when you apply dynamic CoS traffic-shaping attributes to a dynamic interface set or a dynamic ACI interface set and its member subscriber sessions. For complete information about the Junos OS CoS traffic-shaping predefined dynamic variables and RADIUS tag values used with this feature, see "CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets" on page 1005.

- This feature is supported only for dynamically configured and instantiated subscriber interfaces.
- Do not configure a combination of **$junos-cos-*parameter*** predefined dynamic variables with RADIUS tag values in the 100s range and **$junos-cos-*parameter*** predefined dynamic variables with tag values not in the 100s range in the same traffic-control profile. If you do so, the subscriber authentication process fails.
- Use the **$junos-cos-adjust-minimum** predefined variable (tag 109) only in traffic-control profiles for dynamic subscriber interfaces. Using this variable in a traffic-control profile for a dynamic interface set or dynamic ACI interface set has no effect.
- Do not configure the **$junos-cos-excess-rate-high** predefined variable (tag 110) when the **$junos-cos-excess-rate** predefined variable (tag 105) is configured, and vice-versa.
- Do not configure the **$junos-cos-excess-rate-low** predefined variable (tag 111) when the **$junos-cos-excess-rate** predefined variable (tag 105) is configured, and vice-versa.
- Do not configure the **$junos-cos-byte-adjust-frame** predefined variable (tag 114) when the **$junos-cos-byte-adjust** predefined variable (tag 108) is configured, and vice-versa.
- Do not configure the **$junos-cos-byte-adjust-cell** predefined variable (tag 115) when the **$junos-cos-byte-adjust** predefined variable (tag 108) is configured, and vice-versa.
- Use the per-priority **$junos-cos-shaping-rate-*parameter*** predefined variables (tags 116 through 125) only in traffic-control profiles for dynamic interface sets or dynamic ACI interface sets. Using these variables in traffic-control profiles for a dynamic logical subscriber interface causes the subscriber session to fail.

Related Documentation

- Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions on page 1003
- CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets on page 1005
- CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview on page 996

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 55

Configuring RADIUS for Dynamic CoS

- Configuring Initial CoS Parameters Dynamically Obtained from RADIUS on page 999
- Configuring User-Defined CoS Variables in a Dynamic Service Profile on page 1000
- Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions on page 1003
- CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets on page 1005

Configuring Initial CoS Parameters Dynamically Obtained from RADIUS

You can configure a subscriber interface so that subscribers receive initial CoS parameters that the router obtains from the RADIUS authentication server when subscribers log in using that logical interface on the router.

1. Configure external RADIUS server VSAs with values that you expect subscribers to log in with.
   - To configure a RADIUS authentication server to include CoS traffic-shaping parameters in authentication grants on certain subscriber interfaces, configure Juniper Networks VSA 26–108.
   - To configure a RADIUS authentication server to include CoS scheduling and queuing parameters in authentication grants a certain subscriber interfaces, configure Juniper Networks VSA 28–146.

   See "Configuring Router or Switch Interaction with RADIUS Servers" on page 23 and "Configuring RADIUS Server Parameters for Subscriber Access" on page 36.
2. Configure a subscriber interface that supports hierarchical CoS.
   - For static VLAN interfaces, see "Configuring Static Subscriber Interfaces in Dynamic Profiles" on page 765.
   - For static VLAN interfaces over aggregated Ethernet, see "Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet" on page 824.
   - For static IP demux interface sets, see "Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles" on page 771.
   - For dynamic IP demux interface sets, see "Configuring a Subscriber Interface Using a Set of Static IP Demux Interfaces" on page 769

Copyright © 2013, Juniper Networks, Inc.

3. Associate a traffic control profile with the interface.

   See "Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile" on page 979.

4. Configuring initial traffic-shaping parameters to be obtained from RADIUS.

   See "Configuring Dynamic Traffic Shaping and Scheduling Parameters in a Dynamic Profile" on page 972.

5. Configure forwarding classes and scheduler maps statically.

   See *Configuring Forwarding Classes* and *Configuring Scheduler Maps*.

6. Configure a scheduler to specify initial scheduling and queuing parameters to be dynamically obtained from RADIUS when a subscriber logs in.

   See "Configuring Dynamic Schedulers with Variables in a Dynamic Profile" on page 975.

Related Documentation

- Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS on page 989
- Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS on page 1040
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Subscriber Activation and Service Management in an Access Network on page 9
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- Dynamic Profiles Overview on page 638
- Dynamic Variables Overview on page 640
- Junos OS Predefined Variables on page 642

Configuring User-Defined CoS Variables in a Dynamic Service Profile

You can configure user-defined variables in the dynamic service profile for traffic scheduling and shaping parameters.

You can use variables in a dynamic service profile in two ways:

- To enable subscribers to upgrade or downgrade services after login using a RADIUS change of authorization (CoA), configure user-defined variables for CoS parameters as RADIUS attributes.
- To provide subscribers with default values for CoS parameters, configure user-defined variables for CoS parameters with static default values. If you have configured values to be supplied by a RADIUS CoA, subscribers can receive the previously configured default value when deactivating a service.

You activate the variables by referencing them in the traffic control profile configured in the dynamic service profile.

 Copyright © 2013, Juniper Networks, Inc.

To configure user-defined variables for CoS in a dynamic profile:

1. Specify that you want to configure variables in the dynamic profile.

   [edit dynamic-profiles residential-silver variables]

2. Do one of the following to configure variables for the shaping rate:

   - Enable RADIUS to modify the shaping rate based on service changes.

     a. Configure the attribute:

     [edit dynamic-profiles residential-silver variables]
     user@host# set srate radius vendor-id 4874 attribute 108

     b. Configure the tag:

     [edit dynamic-profiles residential-silver variables]
     user@host# set srate radius vendor-id 4874 tag 2

     NOTE: You can configure user-defined values for RADIUS tags that are different than the values that are required in access profiles with predefined variables. For example, in a dynamic service profile, you can assign the shaping rate with a tag of 1 rather than 2, which is required for the $junos-shaping-rate variable. When you configure user-defined values, the VSA that is sent from RADIUS must share the same definition.

   - Configure a default value for the shaping rate.

     [edit dynamic-profiles residential-silver variables]
     user@host# set srate default-value 10m

3. Do one of the following to configure variables for the guaranteed rate:

   - Enable RADIUS to modify the guaranteed rate based on service changes.

     a. Configure the attribute.

     [edit dynamic-profiles residential-silver variables]
     user@host# set grate radius vendor-id 4874 attribute 108

     b. Configure the tag.

     [edit dynamic-profiles residential-silver variables]
     user@host# set grate radius vendor-id 4874 tag 3

   - Configure a default value for the guaranteed rate.

     [edit dynamic-profiles residential-silver variables]
     user@host# set grate default-value 5m

4. Do one of the following to configure variables for the delay buffer rate:

   - Enable RADIUS to modify the delay buffer rate based on service changes.

     a. Configure the attribute.

     [edit dynamic-profiles residential-silver variables]
     user@host# set dbrate radius vendor-id 4874 attribute 108

b. Configure the tag.

```
[edit dynamic-profiles residential-silver variables]
user@host# set dbrate radius vendor-id 4874 tag 4
```

- Configure a default value for the delay buffer rate.

```
[edit dynamic-profiles residential-silver variables]
user@host# set dbrate default-value 10m
```

5. Do one of the following to configure variables for the scheduler map:

- Enable RADIUS to modify the scheduler map based on service changes.

a. Configure the attribute.

```
[edit dynamic-profiles residential-silver variables]
user@host# set smap radius vendor-id 4874 attribute 108
```

b. Configure the tag.

```
[edit dynamic-profiles residential-silver variables]
user@host# set smap radius vendor-id 4874 tag 1
```

- Configure a default value for the scheduler map.

```
[edit dynamic-profiles residential-silver variables]
user@host# set smap default-value triple-play
```

6. Configure the variables for the CoS parameters in the traffic control profile.

Either the shaping rate or the guaranteed rate is required in the traffic control profile.

a. Specify that you want to configure CoS parameters in the dynamic profile.

```
user@host# edit dynamic-profiles residential-silver class-of-service
  traffic-control-profiles tcp1
```

b. Configure the scheduler map variable.

```
[edit dynamic-profiles residential-silver class-of-service traffic-control-profiles
  tcp1]
user@host# set scheduler-map "$smap"
```

c. Configure the shaping rate variable.

```
[edit dynamic-profiles residential-silver class-of-service traffic-control-profiles
  tcp1]
user@host# set shaping-rate "$srate"
```

d. Configure the guaranteed rate variable.

```
[edit dynamic-profiles residential-silver class-of-service traffic-control-profiles
  tcp1]
user@host# set guaranteed-rate "$grate"
```

e. Configure the delay buffer rate variable.

```
[edit dynamic-profiles residential-silver class-of-service traffic-control-profiles
  tcp1]
user@host# set delay-buffer-rate "$dbrate"
```


Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Changing CoS Services Overview on page 993

Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions

To control bandwidth at a household level in a subscriber access network, you can apply RADIUS dynamic class of service (CoS) traffic-shaping attributes to a dynamic interface set or agent-circuit-identifer (ACI) interface set and its member subscriber sessions when the member sessions are authenticated. The dynamic interface set or ACI interface set represents the *household* from which the subscriber sessions originate. The *subscriber sessions*, also referred to as *client sessions* or *subscriber interfaces*, can be dynamic VLAN, PPPoE, or IP demultiplexing (IP demux, for DHCP) subscriber interfaces.

To apply RADIUS dynamic CoS traffic-shaping attributes to both the dynamic interface set and its member subscriber sessions, you must configure two traffic-control profiles in the dynamic profile for the subscriber interface: one traffic-control profile for the "parent" dynamic interface set, and a second traffic-control profile for the dynamic subscriber interfaces. RADIUS tag values for the Junos OS CoS traffic shaping predefined variables used in both traffic-control profiles must be in the 100s range.

Before you begin:

- Create a dynamic profile that defines the VLAN, PPPoE, or IP demux logical subscriber interface.

  See the following topics:

  - Configuring a Basic Dynamic Profile on page 673
  - Configuring VLAN Dynamic Profiles on page 714
  - Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles on page 899
  - Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771

To apply dynamic CoS traffic-shaping attributes to a dynamic ACI or non-ACI interface set and its member subscriber sessions in a dynamic profile for the subscriber interface:

1. Configure two traffic-control profiles at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control profiles]** hierarchy level:

   - Traffic-control profile for the VLAN, PPPoE, or IP demux dynamic subscriber interfaces
   - Traffic-control profile for the dynamic interface set or dynamic ACI interface set to which the subscriber interfaces belong

2. In the traffic-control profiles configured for the dynamic interface set and the subscriber interfaces, reference Junos OS CoS traffic-shaping predefined variables with RADIUS tag values in the 100s range.

See "CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets" on page 1005 for a complete list of the Junos OS predefined variables and RADIUS tag values that you must use in the traffic-control profiles for the dynamic subscriber interfaces and the dynamic interface set.

3. At the **[edit dynamic-profiles *profile-name* interfaces]** hierarchy level, use the **output-traffic-control-profile** statement to apply the traffic-control profiles to the dynamic subscriber interface and the dynamic interface set or dynamic ACI interface set.

Example: Dynamic PPPoE Subscriber Interface over Dynamic ACI Interface Set

The following example shows a dynamic profile named pppoe-subscriber that configures a dynamic PPPoE (**pp0**) subscriber interface over a dynamic ACI interface set.

The **traffic-control-profiles** stanza defines two traffic-control profiles: tcp-pppoe-session for the dynamic PPPoE subscriber interface, and tcp-parent-aci-set for the dynamic "parent" ACI interface set. The **$junos-cos-shaping-rate** predefined variable included in each of these traffic-control profiles is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 102. The **$junos-cos-shaping-mode** variable is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 107.

The **interfaces** stanza applies output traffic-control profile tcp-pppoe-session to the dynamic PPPoE (**pp0**) subscriber interface, and output traffic-control profile tcp-parent-aci-set to the dynamic ACI interface set.

```
[edit dynamic-profiles]
pppoe-subscriber {
  interfaces {
    interface-set "$junos-interface-set-name" {
      interface pp0 {
        unit "$junos-interface-unit";
      }
    }
    pp0 {
      unit "$junos-interface-unit" {
        ppp-options {
          pap;
        }
        pppoe-options {
          underlying-interface "$junos-underlying-interface";
          server;
        }
        no-keepalives;
        family inet {
          unnumbered-address lo0.0;
        }
      }
    }
  }
  class-of-service {
```

 Copyright © 2013, Juniper Networks, Inc.

```
        traffic-control-profiles {
          tcp-pppoe-session {
            scheduler-map smap-1;
            shaping-rate $junos-cos-shaping-rate;
            overhead-accounting $junos-cos-shaping-mode frame-mode-bytes -4
              cell-mode-bytes 12;
          }
          tcp-parent-aci-set {
            shaping-rate $junos-cos-shaping-rate;
            overhead-accounting $junos-cos-shaping-mode frame-mode-bytes -4
              cell-mode-bytes 12;
          }
        }
        interfaces {
          pp0 {
            unit "$junos-interface-unit" {
              output-traffic-control-profile tcp-pppoe-session;
            }
          }
          interface-set $junos-interface-set-name {
            output-traffic-control-profile tcp-parent-aci-set;
          }
        }
      }
      }
    }
```

Related Documentation

- CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets on page 1005
- CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview on page 996
- Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions on page 998

CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets

To control bandwidth at a household level in a subscriber access network, you can apply RADIUS CoS traffic-shaping attributes to a dynamic interface set and its member subscriber sessions when the member sessions are authenticated. The dynamic interface set, which represents the household level in a subscriber access network, can be either a dynamic agent-circuit-identifier (ACI) interface set or a non-ACI–based dynamic interface set. The subscriber sessions belonging to the interface set can be dynamic VLAN, DHCP, or PPPoE subscriber interfaces.

To apply RADIUS CoS traffic-shaping attributes to both the dynamic interface set and its member subscriber sessions, you must configure two traffic-control profiles in the dynamic profile for the subscriber interface: one traffic-control profile for the “parent” dynamic interface set, and a second traffic-control profile for the dynamic subscriber interfaces. RADIUS tag values for the Junos OS CoS traffic-shaping predefined variables used in these traffic-control-profiles must be in the 100s range, as described in Table 84 on page 1006.

To accommodate this feature, the set of existing **$junos-cos-*parameter*** predefined dynamic variables for traffic shaping have been duplicated and assigned a tag value in the 100s range, as listed in Table 84 on page 1006. The tag value is the only difference between the existing predefined dynamic variables and the predefined dynamic variables that you must use with this feature.

For example, the existing **$junos-cos-shaping-rate** predefined variable is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 2. To apply RADIUS CoS traffic-shaping attributes to the dynamic interface set and its member subscriber sessions, you must instead use the **$junos-cos-shaping-rate** predefined variable that is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 102.

Table 84 on page 1006 describes the Junos OS predefined dynamic variables and RADIUS tag values that you can use in a dynamic profile to apply RADIUS CoS traffic-shaping attributes to the dynamic interface set and its member subscriber sessions. The table lists the predefined dynamic variables in ascending order by tag value.

NOTE: All of the predefined variables listed in Table 84 on page 1006 use RADIUS vendor ID 4874 and RADIUS attribute value 108.

Table 84: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets

Predefined Variable	RADIUS Tag Value	Description
$junos-cos-scheduler-map	101	Scheduler-map name configured in a traffic-control profile in a dynamic profile.
$junos-cos-shaping-rate	102	Shaping rate configured in a traffic-control profile in a dynamic profile. Represents the maximum bandwidth of a CoS scheduler node.
$junos-cos-guaranteed-rate	103	Guaranteed rate configured in a traffic-control profile in a dynamic profile. Represents the minimum bandwidth of a CoS scheduler node.
$junos-cos-delay-buffer-rate	104	Delay-buffer rate configured in a traffic-control profile in a dynamic profile.

Copyright © 2013, Juniper Networks, Inc.

Table 84: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets *(continued)*

Predefined Variable	RADIUS Tag Value	Description
$junos-cos-excess-rate	105	Excess rate configured in a traffic-control profile in a dynamic profile; scheduler weighting when operating in the excess region between the guranteed rate and the shaping rate. NOTE: Do not configure the $junos-cos-excess-rate variable when either the $junos-cos-excess-rate-high variable or the $junos-cos-excess-rate-low variable is configured.
$junos-cos-traffic-control-profile	106	Traffic-control profile configured in a dynamic profile for subscriber access.
$junos-cos-shaping-mode	107	Overhead-accounting mode configured in a traffic-control profile in a dynamic profile to shape downstream ATM traffic based on either frames (frame-mode) or cells (cell-mode).
$junos-cos-byte-adjust	108	Byte adjustment value for the cell or frame shaping mode configured in a traffic-control profile in a dynamic profile. NOTE: Do not configure the $junos-cos-byte-adjust variable when either the $junos-cos-byte-adjust-frame variable or the $junos-cos-byte-adjust-cell variable is configured.
$junos-cos-adjust-minimum	109	Minimum adjusted shaping rate configured in a traffic-control profile for a dynamic subscriber interface. Specifying this variable in a traffic-control profile for a dynamic interface set has no effect.

Copyright © 2013, Juniper Networks, Inc.

Table 84: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets *(continued)*

Predefined Variable	RADIUS Tag Value	Description
$junos-cos-excess-rate-high	110	Shaping rate configured for excess high-priority traffic in a traffic-control profile in a dynamic profile. NOTE: Do not configure the $junos-cos-excess-rate-high variable when the $junos-cos-excess-rate variable is configured.
$junos-cos-excess-rate-low	111	Shaping rate configured for excess low-priority traffic in a traffic-control profile in a dynamic profile. NOTE: Do not configure the $junos-cos-excess-rate-low variable when the $junos-cos-excess-rate variable is configured.
$junos-cos-shaping-rate-burst	112	Burst size for the shaping rate configured in a traffic-control profile in a dynamic profile.
$junos-cos-guaranteed-rate-burst	113	Burst size for the guaranteed rate configured in a traffic-control profile in a dynamic profile.
$junos-cos-byte-adjust-frame	114	Overhead bytes when downstream ATM traffic is in frame-mode. NOTE: Do not configure the $junos-cos-byte-adjust-frame variable when the $junos-cos-byte-adjust variable is configured.
$junos-cos-byte-adjust-cell	115	Overhead bytes when downstream ATM traffic is in cell-mode. NOTE: Do not configure the $junos-cos-byte-adjust-cell variable when the $junos-cos-byte-adjust variable is configured.

Copyright © 2013, Juniper Networks, Inc.

Table 84: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets *(continued)*

Predefined Variable	RADIUS Tag Value	Description
$junos-cos-shaping-rate-priority-high	116	Shaping rate configured for high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-priority-high-burst	117	Shaping rate burst size configured for high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-priority-medium	118	Shaping rate configured for medium-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-priority-medium-burst	119	Shaping rate burst size configured for medium-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-priority-low	120	Shaping rate configured for low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.

Copyright © 2013, Juniper Networks, Inc.

Table 84: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets *(continued)*

Predefined Variable	RADIUS Tag Value	Description
$junos-cos-shaping-rate-priority-low-burst	121	Shaping rate burst size configured for low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-excess-high	122	Shaping rate configured for excess high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-excess-high-burst	123	Shaping rate burst size configured for excess high-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-excess-low	124	Shaping rate configured for excess low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.
$junos-cos-shaping-rate-excess-low-burst	125	Shaping rate burst size configured for excess low-priority traffic in a traffic-control profile for a dynamic interface set or dynamic ACI interface set at a household level. Specifying this variable in a traffic-control profile for a dynamic subscriber interface is prohibited.

Related Documentation

- Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions on page 1003
- CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview on page 996

Copyright © 2013, Juniper Networks, Inc.

- Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions on page 998
- Junos OS Predefined Variables on page 642

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 56

Interface Solutions for Dynamic CoS Overview

- CoS for Aggregated Ethernet Subscriber Interfaces Overview on page 1013
- CoS for PPPoE Subscriber Interfaces Overview on page 1014
- CoS for L2TP LAC Subscriber Interfaces Overview on page 1016
- CoS for L2TP LNS Inline Services Overview on page 1018
- CoS for Interface Sets of Subscribers Overview on page 1019

CoS for Aggregated Ethernet Subscriber Interfaces Overview

You can apply static or dynamic hierarchical CoS to a scheduler node at the aggregated Ethernet logical interface, its underlying physical interface, or an interface set.

- Guidelines for Configuring CoS for Aggregated Ethernet Subscribers on page 1013

Guidelines for Configuring CoS for Aggregated Ethernet Subscribers

When you configure CoS for aggregated Ethernet interfaces, consider the following guidelines:

- Configure the aggregated Ethernet logical interface over two physical interfaces capable of performing hierarchical scheduling.
- For VLAN subscriber interfaces over aggregated Ethernet, you must enable link protection on the aggregated Ethernet interface for hierarchical CoS to operate.
- Link protection is not required for IP or demux subscriber interfaces over aggregated Ethernet. We recommend that you enable targeted distribution on the demux interface to provide accurate hierarchical scheduling for these links.
- Keep the following guidelines in mind when configuring interface sets of aggregated Ethernet interfaces:
  - Sets of aggregated Ethernet interfaces are supported on MPC/MIC interfaces on MX Series routers only.
  - The supported logical interfaces for aggregated Ethernet in an interface set include VLAN demux interfaces, IP demux interfaces, and PPPoE logical interfaces over VLAN demux interfaces.

Copyright © 2013, Juniper Networks, Inc.

- The link membership list and scheduler mode of the interface set are inherited from the underlying aggregated Ethernet interface over which the interface set is configured.
- When an aggregated Ethernet interface operates in link protection mode, or if the scheduler mode is configured to replicate member links, the scheduling parameters of the interface set are copied to each of the member links.
- If the scheduler mode of the aggregated Ethernet interface is set to scale member links, the scheduling parameters are scaled based on the number of active member links and applied to each of the aggregated interface member links.

Related Documentation

- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- For hardware requirements, see Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- For configuration instructions, see Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links on page 1023 and Configuring an Interface Set of Subscribers in a Dynamic Profile on page 1028
- For additional information about subscribers over aggregated Ethernet, see Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview on page 816, Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 818, and Static and Dynamic VLAN Subscriber Interfaces over Aggregated Ethernet Overview on page 815.

CoS for PPPoE Subscriber Interfaces Overview

You can configure CoS functionality for static and dynamic PPPoE subscriber interfaces configured on Gigabit Ethernet Intelligent Queuing 2 (IQ2) and Ethernet Enhanced IQ2 (IQ2E) PICs on the M120 and M320 routers, and on MPCs on the MX Series 3D Universal Edge Router.

For all supported hardware platforms, you can attach an output traffic control profile that contains basic shaping and scheduling properties directly to a PPPoE interface. In this type of scenario, you can use each PPPoE interface to represent a household and shape all of the household traffic to an aggregate rate. Each forwarding class is mapped to a queue, and represents one type of services provided to a household customer.

Both the IQ2E PIC and MPC support hierarchical scheduling functionality that is not available on the IQ2 PIC. To shape customer or DSLAM traffic at different levels of the PPPoE interface hierarchy, you can attach traffic control profiles to interface sets that contain PPPoE members.

MPCs support subscriber interfaces with PPPoE encapsulation over aggregated Ethernet interfaces. These PPPoE subscriber interfaces are configured over VLAN demux interfaces, which are also configured over Aggregated Ethernet interfaces.

 Copyright © 2013, Juniper Networks, Inc.

You can configure 802.3ad link aggregation group (LAG) stateful port and dense port concentrator (DPC) redundancy. This provides targeted distribution of non-replicated (stacked) PPPoE or IP dlemux links over VLAN demux links, which in turn are over an aggregated Ethernet (AE) logical interface. Service providers with PPPoE or IP demux interfaces for CoS configurations can provide DPC and port redundancy to subscribers.

NOTE: **For static PPPoE underlying logical interfaces, use PPPoE interface sets.**

Related Documentation

- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 963
- Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 965
- Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface on page 1024
- For more information about the IQ2 and IQ2 PICs, see the *Junos OS Class of Service Library for Routing Devices*

CoS for L2TP LAC Subscriber Interfaces Overview

You can apply CoS to the Layer 2 Tunnel Protocol (L2TP) access concentrator (LAC) component.

In Layer 2 Tunnel Protocol (L2TP) configurations, IP and L2TP headers are added to packets arriving at a PPP subscriber interface on the L2TP access concentrator (LAC) before being tunneled to the L2TP network server (LNS). You can manage the IP header by configuring classifiers and rewrite-rules that transfer the ToS (Type of Service) value or the 802.1p value from the *inner* IP header to the *outer* IP header of the L2TP packet.

Figure 28 on page 1016 shows the classifier and rewrite rules that you can configure from the LAC to the LNS, and from the LNS to the LAC.

Figure 28: CoS Configuration for L2TP LAC Topology

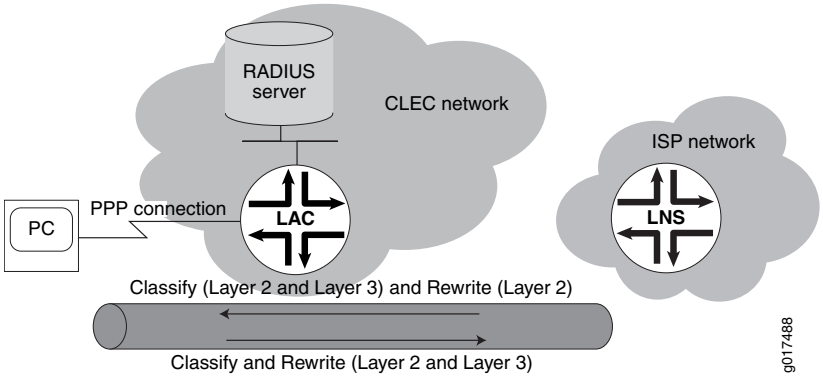

- Traffic from LAC to LNS on page 1016
- LAC Tunnels: Traffic from LNS to LAC on page 1017

Traffic from LAC to LNS

To set the ToS value or the 802.1p value on the inner IP header, you can configure both fixed and behavior aggregate (BA) classifiers for subscribers at Layer 2 or Layer 3 of the network.

Table 85 on page 1016 lists the configuration options for applying classifiers to a subscriber interface on an ingress LAC tunnel.

Table 85: Ingress LAC Tunnel Classifier Options

Classifier	Subscriber Interface
Fixed	Either of the following: • PPP interface • Underlying VLAN interface

Copyright © 2013, Juniper Networks, Inc.

Table 85: Ingress LAC Tunnel Classifier Options *(continued)*

Classifier	Subscriber Interface
Layer 2	Either of the following: • PPP interface • Underlying VLAN interface
Layer 3	Family of PPP interfaces

You cannot configure a Layer 2 and fixed classifier together.

The behavior of the Layer 2 and Layer 3 classifiers depends on the configuration. For example, a Layer 3 classifier for a family of PPP interfaces overrides a Layer 2 classifier configured at the PPP interface, except for the unknown packets and control packets.

If you do not configure a classifier for Layer 2, the system applies the default Layer 3 classifier so that tunneled and terminated subscribers have the same behavior. To prevent unknown packets and control packets from being discarded, the system assigns them to the best-effort forwarding class.

For egress tunnels, you configure rewrite rules at the PPP interface to set the ToS or 802.1p value of the outer IP header. Rewrite rules are applied accordingly to the forwarding class, packet loss priority (PLP), and code point.

LAC Tunnels: Traffic from LNS to LAC

On a LAC, mapping the inner IP header to the outer IP header of the L2TP packet depends on the classifier and rewrite-rule configurations. For example, Table 86 on page 1017 lists the values for the classifier and rewrite rules for a VLAN interface. For assured forwarding, the inner 802.1p value (**ob001**) is classified with the assured-forwarding class and low loss priority at the ingress interface. Based on the assured-forwarding class and low loss priority in the rewrite rule, the ToS value in the outer IP header is set to **ob001**.

Table 86: Sample Result

Inner .1p Value	Forwarding Class	Loss Priority	Code Point	Outer ToS Value
ob000	best-effort	low	000	ob000
ob001	assured-forwarding	low	001	ob001
ob101	expedited-forwarding	low	101	ob101
ob111	network-control	low	11	ob111

Related Documentation

- Configuring Dynamic CoS for an L2TP LAC Tunnel on page 1025

CoS for L2TP LNS Inline Services Overview

You can apply hierarchical scheduling and per-session shaping to Layer 2 Tunnel Protocol (L2TP) network server (LNS) inline services using a static or dynamic CoS configuration.

This feature is supported on MIC and MPC interfaces on MX240, MX480, and MX960 routers.

- Guidelines for Applying CoS to the LNS on page 1018
- Hardware Requirements for Inline Services on the LNS on page 1018

Guidelines for Applying CoS to the LNS

In L2TP configurations, IP, UDP, and L2TP headers are added to packets arriving at a PPP subscriber interface on the L2TP access concentrator (LAC) before being tunneled to the LNS.

When a service interface is configured for an L2TP LNS session, it has an *inner* IP header and an outer IP header. You can configure CoS for an LNS session that corresponds to the inner IP header only. The *outer* IP header is used for L2TP tunnel processing only.

However, we recommend that you configure classifiers and rewrite-rules to transfer the ToS (type of service) value from the inner IP header to the outer IP header of the L2TP packet.

Figure 29 on page 1018 shows the classifier and rewrite rules that you can configure on an LNS inline service.

Figure 29: Processing of CoS Parameters in an L2TP LNS Inline Service

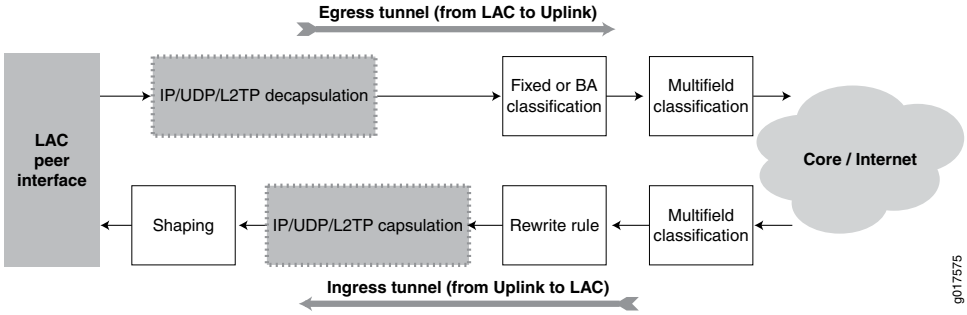

By default, the shaping calculation on the service interface includes the L2TP encapsulation. If necessary, you can configure additional adjustments for downstream ATM traffic from the LAC or differences in Layer 2 protocols.

Hardware Requirements for Inline Services on the LNS

Hierarchical scheduling for L2TP LNS inline services is supported on MIC and MPC interfaces only. The services that you can configure depend on the hardware combination. Table 87 on page 1019 lists the supported inline services and peer interfaces for each MIC and MPC combination.

Copyright © 2013, Juniper Networks, Inc.

Table 87: Hardware Requirements for L2TP LNS Inline Services

MPC Module	Inline Service Support—With Per-Session Shaping	Inline Service Support—Without Per-Session Shaping
MX-MPC1-3D MX-MPC2-3D	No	Yes
MX-MPC1-3D-Q MX-MPC2-3D-Q MX-MPC2-3D-EQ	Yes	Yes
MPC-3D-16XGE-SFPP	No	No

Related Documentation

- *Configuring Static CoS for an L2TP LNS Inline Service*
- Configuring Dynamic CoS for an L2TP LNS Inline Service on page 1026

CoS for Interface Sets of Subscribers Overview

Interface sets enable service providers to group logical interfaces so they can apply CoS parameters to all of the traffic in the group.

Interface sets are beneficial for various scenarios in a subscriber access network. For example, you can use an interface set to configure a local loop with a small number of subscribers. Interface sets are also useful for grouping a large number of subscribers into a particular service class or for defining traffic engineering aggregates for DSLAMs.

- Guidelines for Configuring Dynamic Interface Sets in a Subscriber Access Network on page 1019

Guidelines for Configuring Dynamic Interface Sets in a Subscriber Access Network

Interface sets enable service providers to group logical interfaces so they can apply CoS parameters to all of the traffic in the group.

Interface sets are beneficial for various scenarios in a subscriber access network. For example, you can use an interface set to configure a local loop with a small number of subscribers. Interface sets are also useful for grouping a large number of subscribers into a particular service class or for defining traffic engineering aggregates for DSLAMs.

When configuring interface sets for subscriber access, keep the following guidelines in mind:

- You can configure interface sets of VLAN demux, PPPoE, or demux interfaces over aggregated Ethernet interfaces.
- An interface can only belong to one interface set. If you try to add the same interface to different interface sets, the commit operation fails.

Copyright © 2013, Juniper Networks, Inc.

- You configure the interface set and the traffic scheduling and shaping parameters in a dynamic profile. However, you must apply the traffic-control profile to the interface set in the static **[edit class-of-service]** hierarchy.

  NOTE: **This rule applies to all interface sets except ACI sets.**

- The **$junos-interface-set-name** predefined variable is available only for RADIUS Accept messages; change of authorization (CoA) requests are not supported.
- The **$junos-svlan-interface-set-name** predefined variable locally generates an interface set name for use by dual-tagged VLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is ***physical_interface_name - outer_VLAN_tag***. For example, an aggregated Ethernet interface "ae0," with a dual-tagged VLAN interface that has an outer tag of "111," results in a **$junos-svlan-interface-set-name** dynamic variable of "ae0-111". Similarly, a non-aggregated Ethernet interface of ge-1/1/0, with the same dual-tagged VLAN interface that has an outer tag of "111," results in a **$junos-svlan-interface-set-name** dynamic variable of "ge-1/1/0-111".
- The **$junos-tagged-vlan-interface-set-name** predefined variable locally generates an interface set name used for grouping logical interfaces stacked over logical stacked VLAN demux interfaces for either a 1:1 (dual-tagged; individual client) VLAN or N:1 (single tagged; service) VLAN. The format of the generated variable differs with VLAN type as follows:
  - Dual-tagged (client) VLAN—***physical_interface_name - outer_VLAN_tag - inner_VLAN_tag***. For example, an aggregated Ethernet interface "ae0," with a dual-tagged VLAN interface that has an outer tag of "111" and an inner tag of "200," results in a **$junos-tagged-vlan-interface-set-name** dynamic variable of "ae0-200-111". Similarly, a non-aggregated Ethernet interface of ge-1/1/0, with the same dual-tagged VLAN interface that has an outer tag of "111" and an inner tag of "200," results in a **$junos-tagged-vlan-interface-set-name** dynamic variable of "ge-1/1/0-200-111".
  - Single tagged (service) VLAN—***physical_interface_name - VLAN_tag***. For example, an aggregated Ethernet interface "ae0," with an N:1 VLAN using the single tag of "200," results in a **$junos-tagged-vlan-interface-set-name** dynamic variable of "ae0-200". Similarly, a non-aggregated Ethernet interface of ge-1/1/0, with the same N:1 VLAN using the single tag of "200," results in a **$junos-tagged-vlan-interface-set-name** dynamic variable of "ge-1/1/0-200".
- All dynamic demux, dual-tagged VLAN logical interfaces with the same outer VLAN tag and physical interface are assigned to the same interface set and all CoS values provisioned with the dynamic profile are applied to the interfaces that are part of the set.
- The interface set name must be explicitly referenced in the CoS configuration as part of the static configuration outside of the dynamic profile. The CoS configuration is static and the interface set name must be statically referenced.

 Copyright © 2013, Juniper Networks, Inc.

NOTE: This rule applies to all interface sets except ACI sets.

- RADIUS can return an *access-accept* message under certain conditions. A configured RADIUS VSA for the interface set name takes precedence over the locally generated variable on the router. This means that if the interface-set-name VSA is configured on RADIUS, the router continues to use this variable instead of the locally generated value from the dynamic variable.
- Sets of aggregated Ethernet interfaces are supported on MPC/MIC interfaces on MX Series routers only.
- The supported interface stacks for aggregated Ethernet in an interface set include VLAN demux interfaces, IP demux interfaces, and PPPoE logical interfaces over VLAN demux interfaces.
- The link membership list and scheduler mode of the interface set are inherited from the underlying aggregated Ethernet interface over which the interface set is configured.
- When an aggregated Ethernet interface operates in link protection mode, or if the scheduler mode is configured to replicate member links, the scheduling parameters of the interface set are copied to each of the member links.
- If the scheduler mode of the aggregated Ethernet interface is set to scale member links, the scheduling parameters are scaled based on the number of active member links and applied to each of the aggregated interface member links.

Related Documentation

- Configuring an Interface Set of Subscribers in a Dynamic Profile on page 1028
- Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile on page 1065

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 57

Configuring Interface Solutions for Dynamic CoS

- Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links on page 1023
- Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface on page 1024
- Configuring Dynamic CoS for an L2TP LAC Tunnel on page 1025
- Configuring Dynamic CoS for an L2TP LNS Inline Service on page 1026
- Configuring an Interface Set of Subscribers in a Dynamic Profile on page 1028

Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links

You can enable hierarchical CoS on a subscriber interface with an underlying aggregated Ethernet interface.

Before you begin, configure the subscriber interface with aggregated Ethernet.

- To configure a VLAN interface over aggregated Ethernet with link protection, see "Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet" on page 824 and *Configuring Link Protection for Aggregated Ethernet Interfaces*.
- To configure a demux subscriber interface:

  For static and dynamic IP demux interfaces, see "Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet" on page 825.

  For static and dynamic VLAN demux interfaces, see "Configuring a Static or Dynamic VLAN Demux Subscriber Interface over Aggregated Ethernet" on page 826.

BEST PRACTICE: Link protection is not required for IP or demux subscriber interfaces. We recommend that you enable targeted distribution on the demux interface to provide accurate hierarchical scheduling for these links. See "Providing Accurate Scheduling for a Demux Subscriber Interface of Aggregated Ethernet Links" on page 1095.

To configure hierarchical CoS on the link aggregation (LAG) bundle:

1. Specify that you want to access the LAG bundle.

   user@host# **edit interfaces ae*x***

2. Configure the link aggregation (LAG) bundle with hierarchical scheduler mode.

   [edit interfaces ae*x*]
   user@host# **set hierarchical-scheduler**

You can then attach static or dynamic traffic shaping and scheduling parameters at the aggregated Ethernet logical interface or its underlying physical interface. See:

- Configuring Traffic Scheduling and Shaping for Subscriber Access on page 971
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973
- Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile on page 979

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 986
- CoS for Subscriber Access Overview on page 949
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956

Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface

You can configure hierarchical CoS on a static PPPoE subscriber interface.

Before you begin:

- Configure the static PPPoE subscriber interface.

  See *Configuring PPPoE*.

To configure hierarchical CoS on a static PPPoE subscriber interface:

1. Specify the PPPoE interface that you want to configure.

   user@host# **edit interfaces *pppoe-interface-name***

2. Configure the hierarchical scheduler for the interface.

   [edit interfaces *interface-name*]
   user@host# **set hierarchical-scheduler**

3. (Optional) Group the PPPoE interfaces in an interface set.

   [edit]
   user@host# **edit interfaces interface-set *interface-set-name***

You can now configure static traffic and scheduling parameters for each traffic-control profile, and attach each traffic-control profile to the PPPoE interface or the PPPoE

 Copyright © 2013, Juniper Networks, Inc.

interface set. For more information, see the *Junos OS Class of Service Library for Routing Devices*.

Related Documentation

- For hardware requirements and configuration guidelines, see Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for PPPoE Subscriber Interfaces Overview on page 1014
- Example: Configuring Hierarchical Scheduling and Queuing for a Static PPPoE Subscriber Interface on page 1045
- Example: Configuring Hierarchical Scheduling and Queuing for an Underlying Static PPPoE Subscriber Interface on page 1047
- Example: Configuring Hierarchical Scheduling and Queuing for an Interface Set of Static PPPoE Subscriber Interfaces on page 1050
- Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 986
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956

Configuring Dynamic CoS for an L2TP LAC Tunnel

In L2TP configurations, IP and L2TP headers are added to packets arriving at a PPP subscriber interface on the LAC before being tunneled to the L2TP network server (LNS).

Classifiers and rewrite rules enable you to properly transfer the ToS (Type of Service) value or the 802.1p value from the inner IP header to the outer IP header of the L2TP packet.

Before you begin, configure the L2TP LAC. See "Configuring an L2TP LAC" on page 402.

To manage the IP header values for a LAC tunnel:

1. Configure the classifier for the inner tunnel.
   a. Define the fixed or behavior aggregate (BA) classifier.
      - To configure a fixed classifier:

        [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]
        user@host# **set forwarding-class *class-name***

      - To configure a BA classifier:

        [edit class-of-service]
        user@host#**set classifiers (ieee-802.1 | inet-precedence) *classifier-name* forwarding-class *class-name* loss-priority *level* code-points [*aliases*] [*bit-patterns*]**

   b. Apply the classifier to the Layer 2 interface or Layer 3 interface. For Layer 2, you can apply the classifier at the PPP interface or an underlying VLAN interface. For Layer 3, you can apply classifiers to a family of PPP interfaces.
      - To apply the classifier for the IEEE 802.1p value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
  unit logical-unit-number classifiers]
user@host# set ieee-802.1 (classifier-name | default) vlan-tag (inner | outer)
```

- To apply the classifier for the ToS value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
  unit logical-unit-number classifiers]
user@host# set inet-precedence (classifier-name | default)
```

2. Configure the rewrite rule for the egress tunnel.

   a. Configure the rewrite rule with the forwarding class and the loss priority value.

```
[edit class-of-service]
user@host# set rewrite-rules (ieee-802.1 | inet-precedence) rewrite-name
  forwarding-class class-name loss-priority level code-point (alias | bits)
```

   b. Apply the rewrite rule to the PPP interface for which the L2TP tunnel is configured.

   - To apply the rewrite-rule for the IEEE 802.1p value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
  unit logical-unit-number rewrite-rules]
user@host# set ieee-802.1 (rewrite-name | default) vlan-tag (outer |
  outer-and-inner)
```

   - To apply the rewrite rule for the ToS value:

```
[edit dynamic-profiles profile-name class-of-service interfaces interface-name
  unit logical-unit-number rewrite-rules]
user@host# set inet-precedence (rewrite-name | default)
```

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for L2TP LAC Subscriber Interfaces Overview on page 1016

Configuring Dynamic CoS for an L2TP LNS Inline Service

You can configure hierarchical scheduling for an L2TP LNS inline service and manage the IP header values using rewrite rules and classifiers.

Before you begin, configure the L2TP LNS inline service interface. See "Configuring an L2TP LNS with Inline Service Interfaces" on page 417.

To configure CoS for an L2TP LNS inline service in a dynamic profile:

1. Configure the hierarchical scheduler for the service interface (**si**) interface.

```
[edit interfaces si-fpc/port/pic ]
user@host# set hierarchical-scheduler maximum-hierarchy-levels 2
```

BEST PRACTICE: To enable Level 3 nodes in the LNS scheduler hierarchy and to provide better scaling, we recommend that you also specify a maximum of two hierarchy levels.

Copyright © 2013, Juniper Networks, Inc.

2. Configure the LNS to reflect the IP ToS value in the inner IP header to the outer IP header.

   [edit services l2tp tunnel-group *name*]
   user@host# set tos-reflect

3. Configure the classifier for egress traffic from the LAC.

   a. Define the fixed or behavior aggregate (BA) classifier.

      - To configure a fixed classifier:

        [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]
        user@host# set forwarding-class *class-name*

      - To configure a BA classifier:

        [edit class-of-service]
        user@host# set classifiers (dscp | dscp-ipv6 | inet-precedence) *classifier-name* forwarding-class *class-name* loss-priority *level* code-points [*aliases*] [*bit-patterns*]

   b. Apply the classifier to the service interface.

      - To apply the classifier for the DSCP or DSCP IPv6 value:

        [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers]
        user@host# set dscp (*classifier-name* | default)
        user@host# set dscp-ipv6 (*classifier-name* | default)

      - To apply the classifier for the ToS value:

        [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers]
        user@host# set inet-precedence (*classifier-name* | default)

4. Configure and apply a rewrite-rule to ingress traffic to the LAC:

   a. Configure the rewrite rule with the forwarding class and the loss priority value.

      [edit class-of-service]
      user@host# set rewrite-rules (dscp | dscp-ipv6 | inet-precedence) *rewrite-name* forwarding-class *class-name* loss-priority *level* code-point (*alias* | *bits*)

   b. Apply the rewrite rule to the service interface.

      - To apply the rewrite rule for the DSCP or DSCP IPv6 value:

        [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]
        user@host# setdscp (*rewrite-name* | default)
        user@host# set dscp-ipv6 (*rewrite-name* | default)

      - To apply the rewrite rule for the ToS value:

        [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]
        user@host# set inet-precedence (*rewrite-name* | default)

5. (Optional) Configure additional adjustments for downstream ATM traffic.

By default, the shaping calculation on the service interface includes the L2TP encapsulation. If necessary, you can configure additional adjustments for downstream ATM traffic from the LAC or differences in Layer 2 protocols.

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles profile-name]
user@host# set overhead-accounting (frame-mode | cell-mode |
  $junos-cos-shaping-mode) <bytes (byte-value | $junos-cos-byte-adjust)
```

6. Apply the traffic-control profile.

```
[edit dynamic-profiles profile-name class-of-service interfaces
  $junos-interface-ifd-name unit $junos-interface-unit]
user@host# set output-traffic-control-profile profile-name
```

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for L2TP LNS Inline Services Overview on page 1018
- Example: Configuring an L2TP LNS on page 440
- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 1092

Configuring an Interface Set of Subscribers in a Dynamic Profile

Interface sets enable you to provide hierarchical scheduling to a group of subscriber interfaces.

Before you begin, configure the subscriber interfaces that you intend to include in the interface set.

- For static VLAN interfaces, see "Configuring Static Subscriber Interfaces in Dynamic Profiles" on page 765.
- For dynamic VLAN interfaces, see "Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet" on page 824.
- For dynamic IP demux interfaces, see "Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles" on page 771 and "Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet" on page 825.
- For dynamic VLAN demux interfaces, see "Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles" on page 772.
- For dynamic PPPoE interfaces, see "Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles" on page 899.
- For aggregated Ethernet interfaces, see "Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links" on page 1023

To configure an interface set of subscriber interfaces:

1. Configure the interface set in the dynamic profile.

```
[edit dynamic-profiles profile-name interfaces]
user@host#edit interface-set interface-set-name
```

 Copyright © 2013, Juniper Networks, Inc.

Replacing the ***interface-set-name*** variable with the **$junos-interface-set-name**, **$junos-svlan-interface-set-name**, or **$junos-tagged-vlan-interface-set-name** predefined variable. The interface set is created dynamically when the subscriber logs in.

2. Include the interfaces within the dynamic interface-set.

   [edit dynamic-profiles *profile-name* interfaces interface-set $junos-interface-set-name]
   user@host# **set interface *interface-name* unit *logical-unit-number***

3. Apply traffic shaping and queuing parameters to the interface set.

   TIP: You must configure the interface set in the static **[edit class-of-service]** hierarchy, not in the **[edit dynamic-profiles]** hierarchy.

   [edit class-of-service interfaces]
   user@host# **edit interface-set *interface-set-name***
   [edit class-of-service interfaces interface-set *interface-set-name*]
   user@host# **set output-traffic-control-profile *profile-name***

Related Documentation

- CoS for Interface Sets of Subscribers Overview on page 1019
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- CoS for Interface Sets of Subscribers Overview on page 1019
- Example: Configuring a Dynamic Interface Set of VLAN Subscribers on page 1052
- CoS for Aggregated Ethernet Subscriber Interfaces Overview on page 1013

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 58

Dynamic CoS for Subscriber Access Examples

- Example: Configuring Static Hierarchical Scheduling and Queuing for Subscriber Access on page 1031
- Example: Configuring Dynamic Hierarchical Scheduling and Queuing for Subscriber Access on page 1033
- Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS on page 1040
- Example: Providing Unique Rate Configurations for Schedulers in a Dynamic Profile on page 1042
- Example: Configuring Aggregate Scheduling of Queues for Residential Subscribers on Static IP Demux Interfaces on page 1043
- Example: Configuring Hierarchical Scheduling and Queuing for a Static PPPoE Subscriber Interface on page 1045
- Example: Configuring Hierarchical Scheduling and Queuing for an Underlying Static PPPoE Subscriber Interface on page 1047
- Example: Configuring Hierarchical Scheduling and Queuing for an Interface Set of Static PPPoE Subscriber Interfaces on page 1050
- Example: Configuring a Dynamic Interface Set of VLAN Subscribers on page 1052
- Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile on page 1065

Example: Configuring Static Hierarchical Scheduling and Queuing for Subscriber Access

This example shows you how to configure CoS for a subscriber in a dynamic profile. The CoS parameters configure a best-effort, data service for subscribers.

1. Configure the static CoS parameters in the **[edit class-of-service]** hierarchy.

   You must configure the scheduler maps in this hierarchy; it will get referenced in the dynamic profile.

```
class-of-service {
  forwarding-classes {
    queue 0 best-effort;
    queue 1 expedited-forwarding;
```

```
        queue 3 network-control;
        queue 2 assured-forwarding;
      }
      scheduler-maps {
        data_smap {
          forwarding-class best-effort scheduler be_sch;
        }
      }
      schedulers {
        be_sch {
          transmit-rate percent 10;
          buffer-size remainder;
          priority low;
        }
      }
    }
```

2. Configure the subscriber interface in the **[edit interfaces]** hierarchy. Enable hierarchical scheduling for the interface.

```
interfaces {
  ge-2/2/0 {
    hierarchical-scheduler;
    vlan-tagging;
    unit 100 {
      vlan-id 100;
      family inet {
        unnumbered-address lo0.0 preferred-source-address 100.0.0.1;
      }
    }
  }
}
```

3. Configure CoS in the dynamic profile.

```
dynamic-profiles {
  data-service {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit" {
          family inet;
        }
      }
    }
    class-of-service {
      traffic-control-profiles {
        tcp1 {
          scheduler-map data_smap;
          shaping-rate 50k;
          guaranteed-rate 10k;
        }
      }
      interfaces {
        "$junos-interface-ifd-name" {
          unit "$junos-underlying-interface-unit" {
            output-traffic-control-profile tcp1;
          }
```

 Copyright © 2013, Juniper Networks, Inc.

```
            }
          }
        }
      }
    }
```

Related Documentation

- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Changing CoS Services Overview on page 993
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950

Example: Configuring Dynamic Hierarchical Scheduling and Queuing for Subscriber Access

In this example, subscribers are provided with a data and voice service defined in an access profile when they initially log in. The RADIUS administrator supplies the initial values on the RADIUS server, and the service activation is performed at subscriber login.

After the initial login, the subscriber adds an assured forwarding service that is not defined in the original access profile. A service profile is used to configure the schedulers and a RADIUS CoA activates the service. The queues defined for the schedulers in the initial scheduler map and the new scheduler map are merged.

In addition, the values for the initial data and voice service are upgraded by the RADIUS administrator through a separate RADIUS CoA message.

To configure the initial service and enable the activation through a RADIUS CoA:

1. Configure the access profile for the service activation.

   a. Configure the VLAN interface for the access profile.

   ```
   [edit]
   dynamic-profiles access-profile {
     interfaces {
       $junos-interface-ifd-name {
         unit $junos-underlying-interface-unit {
           family inet;
         }
       }
     }
   }
   ```

   b. Configure the class of service parameters in the access profile. In this example, you configure Junos OS predefined variables that provide the initial scheduler name and scheduler parameters obtained from the RADIUS authentication server when the subscriber logs in.

      Include the configurations for the interfaces, schedulers, and the scheduler maps.

   ```
   [edit]
   dynamic-profiles access-profile {
     class-of-service {
   ```

```
        traffic-control-profiles {
          tcp1 {
            scheduler-map $junos-cos-scheduler-map;
            shaping-rate $junos-cos-shaping-rate;
            guaranteed-rate $junos-cos-guaranteed-rate;
            delay-buffer-rate $junos-cos-delay-buffer-rate;
          }
        }
        interfaces {
          $junos-interface-ifd-name {
            unit "$junos-underlying-interface-unit" {
              classifiers {
                ieee-802.1 l2_classifier;
              }
              rewrite-rules {
                ieee-802.1 l2_rewrite;
              }
            output-traffic-control-profile tcp1;
            }
          }
        }
        schedulers {
          $junos-cos-scheduler {
            buffer-size percent $junos-cos-scheduler-bs;
            priority $junos-cos-scheduler-pri;
            transmit-rate percent $junos-cos-scheduler-tx;
            drop-profile-map loss-priority low protocol any $junos-cos-scheduler-low;
            drop-profile-map loss-priority medium-low protocol any
              $junos-cos-scheduler-medium-low;
            drop-profile-map loss-priority medium-high protocol any
              $junos-cos-scheduler-medium-high;
            drop-profile-map loss-priority high protocol any $junos-cos-scheduler-high;
          }
        }
        scheduler-maps {
          data_voice_smap {
            forwarding-class be scheduler be_sch;
            forwarding-class ef scheduler ef_sch;
          }
        }
      }
    }
```

Table 88 on page 1034 lists the initial values defined by the RADIUS administrator for the scheduler map and shaping rates.

Table 88: Initial Scheduler Map and Shaping Values at Subscriber Login

Predefined Variable	RADIUS Tag	Value
$junos-cos-scheduler-map	T01	data_voice_smap
$junos-cos-shaping-rate	T02	6m
$junos-cos-guaranteed-rate	T03	4m

Copyright © 2013, Juniper Networks, Inc.

Table 88: Initial Scheduler Map and Shaping Values at Subscriber Login *(continued)*

Predefined Variable	RADIUS Tag	Value
$junos-cos-delay-buffer-rate	T04	4m

Table 89 on page 1035 lists the initial values defined by the RADIUS administrator for the voice (expedited forwarding) scheduler.

Table 89: Initial CoS Values for the Voice Scheduler at Subscriber Login

Predefined Variable	Tag	Value
$junos-cos-scheduler	—	ef_sch
$junos-cos-scheduler-tx	T01	10
$junos-cos-scheduler-bs	T02	10
$junos-cos-scheduler-pri	T03	medium-high
$junos-cos-scheduler-dropfile-low	T04	d3
$junos-cos-scheduler-dropfile-medium-low	T05	d2
$junos-cos-scheduler-dropfile-medium-high	T06	d1
$junos-cos-scheduler-dropfile-high	T07	d0

Table 90 on page 1035 lists the initial values defined by the RADIUS administrator for the data (best effort) scheduler.

Table 90: Initial CoS Values for the Data Scheduler at Subscriber Login

Predefined Variable	Tag	Value
$junos-cos-scheduler	—	be_sch
$junos-cos-scheduler-tx	T01	10
$junos-cos-scheduler-bs	T02	10
$junos-cos-scheduler-pri	T03	low
$junos-cos-scheduler-dropfile-low	T04	d0
$junos-cos-scheduler-dropfile-medium-low	T05	d1
$junos-cos-scheduler-dropfile-medium-high	T06	d2
$junos-cos-scheduler-dropfile-high	T07	d3

Copyright © 2013, Juniper Networks, Inc.

2. Configure the classifiers, drop profiles, forwarding classes, and rewrite rules in the static **[edit class-of-service]** hierarchy.

```
[edit]
class-of-service {
  classifiers {
    dscp dscp_classifier {
      forwarding-class be {
        loss-priority low code-points 000000;
      }
      forwarding-class af {
        loss-priority medium-low code-points 000001;
      }
    }
    ieee-802.1 l2_classifier {
      forwarding-class be {
        loss-priority medium-low code-points 000;
      }
      forwarding-class ef {
        loss-priority medium-low code-points 100;
      }
      forwarding-class af {
        loss-priority medium-low code-points 010;
      }
    }
  }
  drop-profiles {
    d0 {
      fill-level 25 drop-probability 100;
      fill-level 0 drop-probability 0;
    }
    d1 {
      fill-level 50 drop-probability 100;
      fill-level 0 drop-probability 0;
    }
    d2 {
      fill-level 75 drop-probability 100;
      fill-level 0 drop-probability 0;
    }
    d3 {
      fill-level 0 drop-probability 0;
      fill-level 100 drop-probability 100;
    }
  }
  forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
  }
  interfaces {
    ge-1/2/9 {
      shaping-rate 100m;
    }
  }
  rewrite-rules {
```

 Copyright © 2013, Juniper Networks, Inc.

```
        ieee-802.1 l2_rewrite {
          forwarding-class be {
            loss-priority medium-low code-point 000;
          }
          forwarding-class ef {
            loss-priority medium-low code-point 001;
          }
          forwarding-class af {
            loss-priority medium-low code-point 100;
          }
        dscp l2_rewrite {
          forwarding-class be {
            loss-priority medium-low code-points 000;
          }
          forwarding-class ef {
            loss-priority medium-low code-points 001;
          }
          forwarding-class af {
            loss-priority medium-low code-points 001;
          }
      }
    }
```

3. Configure the service profile enable RADIUS to activate the video service after login. The video service corresponds to assured forwarding PHB.

   In this example, you configure Junos OS predefined variables that provide the initial scheduler name and scheduler parameters obtained from the RADIUS authentication server when the subscriber logs in.

```
[edit]
dynamic-profiles service-af {
  variables {
    af_fc default-value video;
    af_sch default-value af_sch;
    sch-drop-any default-value all;
    sch-pri-2 default-value strict-high;
    sch-bs-2 default-value 40;
    sch-tx-2 default-value 3m;
    smap default-value any
  }
    class-of-service {
      scheduler-maps {
        "$smap" {
          forwarding-class “$af_fc” scheduler “$af_sch”;
        }
      }
      schedulers {
        "$af_sch" {
          transmit-rate percent "$sch-tx-2";
          buffer-size percent "$sch-bs-2";
          priority "$sch-pri-2";
          drop-profile-map loss-priority any protocol any drop-profile “$sch-drop-any”;
        }
      }
    }
```

```
}
```

After the three services are activated, subscribers receive upgraded values for the data and voice service when RADIUS sends a change of authorization (CoA). In this case, the CoS parameters are replaced, because multiple subscribers were not enabled on the logical interface.

Table 91 on page 1038 lists the upgraded values defined by the RADIUS administrator.

Table 91: Upgraded CoS Values for the Video Service

Variable	RADIUS Tag	Value
junos-cos-scheduler-map	T01	data_voice_smap
junos-cos-shaping-rate	T02	14m
junos-cos-guaranteed-rate	T03	13m
junos-cos-delay-buffer-rate	T04	12m

Table 92 on page 1038 lists the values defined by the RADIUS administrator for the video (assured forwarding) scheduler.

Table 92: Upgraded CoS Values for the Video Scheduler

Predefined Variable	Tag	Value
$junos-cos-scheduler	—	af_sch
$junos-cos-scheduler-tx	T01	10
$junos-cos-scheduler-bs	T02	10
$junos-cos-scheduler-pri	T03	medium
$junos-cos-scheduler-dropfile-low	T04	d3
$junos-cos-scheduler-dropfile-medium-low	T05	d2
$junos-cos-scheduler-dropfile-medium-high	T06	d1
$junos-cos-scheduler-dropfile-high	T07	d0

Table 93 on page 1039 lists the values defined by the RADIUS administrator for the expedited forwarding scheduler in the CoA message. The values are the same as the initial service.

Copyright © 2013, Juniper Networks, Inc.

Table 93: Initial CoS Values for the Expedited Forwarding Scheduler at Subscriber Login

Predefined Variable	Tag	Value
$junos-cos-scheduler	—	ef_sch
$junos-cos-scheduler-tx	T01	10
$junos-cos-scheduler-bs	T02	10
$junos-cos-scheduler-pri	T03	medium-high
$junos-cos-scheduler-dropfile-low	T04	d3
$junos-cos-scheduler-dropfile-medium-low	T05	d2
$junos-cos-scheduler-dropfile-medium-high	T06	d1
$junos-cos-scheduler-dropfile-high	T07	d0

Table 94 on page 1039 lists the values defined by the RADIUS administrator for the best effort scheduler in the CoA message. The values are the same as the initial service.

Table 94: Initial CoS Values for the Best Effort Scheduler at Subscriber Login

Predefined Variable	Tag	Value
$junos-cos-scheduler	—	be_sch
$junos-cos-scheduler-tx	T01	10
$junos-cos-scheduler-bs	T02	10
$junos-cos-scheduler-pri	T03	low
$junos-cos-scheduler-dropfile-low	T04	d0
$junos-cos-scheduler-dropfile-medium-low	T05	d1
$junos-cos-scheduler-dropfile-medium-high	T06	d2
$junos-cos-scheduler-dropfile-high	T07	d3

Related Documentation

- Changing CoS Services Overview on page 993
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS

The following configuration is an example of a client dynamic profile in which initial CoS parameters are dynamically obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is applied.

For this example, assume that the RADIUS authentication server has been configured with traffic-shaping parameters (at Juniper Networks VSA 26-108) and CoS scheduling and queuing parameters (at Juniper Networks VSA 26–146).

The subscriber interface is a single-unit static gigabit Ethernet VLAN interface on an EQ DPC port:

```
[edit]
interfaces {
  ge-9/0/3 {
    hierarchical-scheduler;
    vlan-tagging;
    unit 100 {
      vlan-id 100;
      family inet {
        address 192.168.32.2/24;
      }
    }
  }
}
```

The client dynamic profile **residential_silver** attaches the traffic-control profile **tcp_1** to the subscriber interface that is defined in the dynamic profile using the **$junos-interface-ifd-name** predefined variable.

```
[edit]
dynamic-profiles {
  residential_silver {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit" {
          family inet;
        }
      }
    }
    class-of-service {
      interfaces {
        "$junos-interface-ifd-name" {
          unit "$junos-underlying-interface-unit" {
            output-traffic-control-profile tcp_1;
          }
        }
      }
    }
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

The traffic-control profile **tcp_1**, references Junos OS predefined variables to obtain a scheduler-map name and traffic-shaping parameter values from RADIUS when a subscriber logs in. For this example, assume that the RADIUS server replaces the Junos OS predefined variable **$junos-cos-scheduler-map** scheduler-map name **business_smap_1**. The scheduler map **business_smap_1** is configured in the client dynamic profile:

```
[edit]
dynamic-profiles {
  residential_silver {
    class-of-service {
      traffic-control-profiles {
        tcp_1 {
          scheduler-map “$junos-cos-scheduler-map”; # ’business_smap_1’
          shaping-rate "$junos-cos-shaping-rate";
          guaranteed-rate "$junos-cos-guaranteed-rate";
          delay-buffer-rate "$junos-cos-delay-buffer-rate";
        }
      }
      scheduler-maps {
        business_smap_1 {
          forwarding-class best-effort scheduler be_sched;
          forwarding-class ef scheduler home_sched
        }
      }
    }
  }
}
```

A scheduler definition references Junos OS predefined variables to obtain scheduler configurations from RADIUS when a subscriber logs in. For this example, assume that the RADIUS server provides scheduler configurations for schedulers named **be_sched** and **home_sched**, which are included in the scheduler map **business_smap_1**:

```
[edit]
dynamic-profiles {
  residential_silver {
    class-of-service {
      schedulers {
        “$junos-cos-scheduler” { # ’be_sched’ and ’home_sched’
          transmit-rate "$junos-cos-scheduler-tx";
          buffer-size "$junos-cos-scheduler-bs";
          priority "$junos-cos-scheduler-pri";
          drop-profile-map loss-priority low protocol any drop-profile
            “$junos-cos-scheduler-dropfile-low“;
          drop-profile-map loss-priority medium-low protocol any drop-profile
            “$junos-cos-scheduler-dropfile-medium-low“;
          drop-profile-map loss-priority medium-high protocol any drop-profile
            “$junos-cos-scheduler-dropfile-medium-high“;
          drop-profile-map loss-priority high protocol any drop-profile
            “$junos-cos-scheduler-dropfile-high“;
        }
      }
    }
  }
}
```

Static configurations for CoS consist of configurations for the forwarding classes used in the scheduler map **business_smap_1** and configurations for drop-profile names provided by RADIUS for as part of the scheduler configurations provided (for **be_sched** and **home_sched**) when a subscriber logs in:

```
[edit]
  class-of-service {
    forwarding-classes {
      queue 0 best-effort;
      queue 1 ef;
    }
    drop-profiles {
      . . . configurations_for_drop_profile_names_provided_by_RADIUS . . .
    }
  }
}
```

Related Documentation

- Subscriber Activation and Service Management in an Access Network on page 9
- Dynamic Profiles Overview on page 638
- Dynamic Variables Overview on page 640
- Junos OS Predefined Variables on page 642
- Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS on page 989
- Configuring Initial CoS Parameters Dynamically Obtained from RADIUS on page 999

Example: Providing Unique Rate Configurations for Schedulers in a Dynamic Profile

Combining static and dynamic schedulers in a dynamic profile enables you to provide subscribers with services that have unique scheduler definitions.

In this example, the network administrator configures the data service with a **transmit-rate** that is rate controlled using the **$junos-cos-scheduler-tx** predefined variable. RADIUS dynamically supplies the percentage value for the transmission rate that is specified in the RADIUS VSA to the data scheduler when the subscriber logs in.

For the best-effort service, the network administrator assigns the remaining transmission rate that is available.

```
schedulers {
  data-scheduler {
    transmit-rate percent rate-limit $junos-cos-scheduler-tx;
    buffer-size percent $junos-cos-scheduler-bs;
    priority $junos-cos-scheduler-pri;
    drop-profile-map loss-priority low protocol any drop-profile d0;
    drop-profile-map loss-priority medium-low protocol any drop-profile d1;
    drop-profile-map loss-priority medium-high protocol any drop-profile d2;
    drop-profile-map loss-priority high protocol any drop-profile d3;
    drop-profile-map loss-priority any protocol any drop-profile all;
  }
  best-effort-scheduler {
```

 Copyright © 2013, Juniper Networks, Inc.

```
    transmit-rate remainder;
    buffer-size percent $junos-cos-scheduler-bs;
    priority medium-high;
    drop-profile-map loss-priority low protocol any drop-profile
      $junos-cos-scheduler-dropfile-low;
    drop-profile-map loss-priority medium-low protocol any drop-profile d1;
    drop-profile-map loss-priority medium-high protocol any drop-profile
      $junos-cos-scheduler-dropfile-medium-high;
    drop-profile-map loss-priority high protocol any drop-profile d3;
    drop-profile-map loss-priority any protocol any drop-profile
      $junos-cos-scheduler-dropfile-any;
  }
```

Related Documentation

- Configuring a Combination of Static and Dynamic Scheduler Parameters in a Scheduler Definition on page 976

Example: Configuring Aggregate Scheduling of Queues for Residential Subscribers on Static IP Demux Interfaces

In this example, scheduling is configured for a residential subscriber. Each forwarding class represents a multiplay service (voice, video, and data), and is equivalent to a queue.

An interface set of IP demux interfaces represents a DSLAM, and provides shaping of subscribers services to a DSLAM aggregate rate.

```
[edit]
interfaces {
  interface-set demux-set {
    interface demux0 {
      unit 0;
      unit 1;
    }
  }
  ge-2/0/1 {
    vlan-tagging;
    unit 1 {
      per-session-scheduler;
      vlan-id 1;
      demux-source inet;
      family inet {
        address 4.4.4.4/24;
      }
    }
  }
  demux0 {
    unit 0 {
      demux-options {
        underlying-interface ge-2/0/1.1;
      }
      family inet {
        address 1.1.1.1/24;
        demux-source {
          1.1.1.0/24;
        }
```

```
            }
        }
        unit 1 {
            demux-options {
                underlying-interface ge-2/0/1.1;
            }
            family inet {
                address 1.1.2.1/24;
                demux-source {
                    1.1.2.0/24;
                }
            }
        }
    }
}
class-of-service {
    traffic-control-profiles {
        T1 {
            scheduler-map m1;
            shaping-rate 5m;
        }
        T2 {
            shaping-rate 60m;
        }
    }
    interfaces {
        interface-set demux-set {
            output-traffic-control-profile T2;
        }
        demux0 {
            unit 0 {
                output-traffic-control-profile T1;
            }
            unit 1 {
                output-traffic-control-profile T1;
            }
        }
    }
    scheduler-maps {
        m1 {
            forwarding-class best-effort scheduler s0;
            forwarding-class expedited-forwarding scheduler s1;
            forwarding-class assured-forwarding scheduler s2;
            forwarding-class network-control scheduler s3;
        }
    }
    schedulers {
        s0 {
            transmit-rate percent 10;
            buffer-size percent 10;
        }
        s1 {
            transmit-rate percent 20;
            buffer-size percent 20;
        }
        s2 {
```

 Copyright © 2013, Juniper Networks, Inc.

```
            transmit-rate percent 30;
            buffer-size percent 30;
        }
        s3 {
            transmit-rate percent 40;
            buffer-size percent 40;
        }
    }
}
```

Related Documentation

- *CoS and Static IP Demux Interface Set Overview*
- Configuring Static Subscriber Interfaces Using IP Demux Interfaces on page 766

Example: Configuring Hierarchical Scheduling and Queuing for a Static PPPoE Subscriber Interface

In this example, the network administrator defines hierarchical queuing and scheduler parameters by configuring traffic control profile and binding it directly to a PPPoE subscriber interface.

This configuration is supported on the IQ2E PIC.

To use this configuration in a broadband access network, each forwarding class can represent one type of services provided to a household customer and is mapped to a queue. Each PPPoE interface represents a household and provides shaping of all household traffic to an aggregate rate. All of the PPPoE interfaces on the physical interfaces are shaped to the underlying physical interface rate.

Table 95 on page 1045 lists the scheduler and queue mapping for this configuration.

Table 95: Scheduler Per Logical Interface Mapping

Level	Type	Mapping
4	Queue	PPPoE interface
3	Scheduler	PPPoE interface
2	Scheduler	—
1	Scheduler	Underlying physical interface

```
interfaces {
    ge-3/0/3 {
        hierarchical-scheduler;
        vlan-tagging;
        unit 0 {
            encapsulation ppp-over-ether;
            vlan-id 100;
        }
    }
    pp0 {
```

```
        unit 0 {
          pppoe-options {
            underlying-interface ge-3/0/3.0;
            server;
          }
          family inet {
            address 120.20.20.20/32 {
              destination 120.20.20.21;
            }
          }
        }
        unit 1 {
          pppoe-options {
            underlying-interface ge-3/0/3.0;
            server;
          }
          family inet {
            address 130.30.30.30/32 {
              destination 130.30.30.31;
            }
          }
        }
        unit 2 {
          pppoe-options {
            underlying-interface ge-3/0/3.0;
            server;
          }
          family inet {
            address 140.40.40.40/32 {
              destination 140.40.40.41;
            }
          }
        }
      }
    }
    class-of-service {
      traffic-control-profiles {
        tcp {
          scheduler-map data_smap;
          shaping-rate 50k;
          guaranteed-rate 10k;
        }
      }
      interfaces {
        pp0 {
          unit 0 {
            output-traffic-control-profile tcp;
          }
            unit 1 {
              output-traffic-control-profile tcp;
              unit 2 {
                output-traffic-control-profile tcp;
                }
                }
              }
              forwarding-classes {
```

 Copyright © 2013, Juniper Networks, Inc.

```
    queue 0 be;
    queue 1 ef;
    queue 3 nc;
    queue 2 af;
  }
  scheduler-maps {
    data_smap {
      forwarding-class be scheduler be_sch;
    }
    voice_data_smap {
      forwarding-class be scheduler be_sch;
    }
    vid_data_smap {
      forwarding-class ef scheduler ef_sch;
    }
  }
  schedulers {
    be_sch {
      transmit-rate percent 10;
      buffer-size remainder;
      priority low;
    }
    ef_sch {
      transmit-rate percent 10;
      buffer-size remainder;
      priority low;
    }
    af_sch {
      transmit-rate percent 10;
      buffer-size remainder;
      priority low;
    }
    nc_sch {
      transmit-rate percent 10;
      buffer-size remainder;
      priority low;
    }
  }
```

Related Documentation

- CoS for PPPoE Subscriber Interfaces Overview on page 1014
- Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface on page 1024

Example: Configuring Hierarchical Scheduling and Queuing for an Underlying Static PPPoE Subscriber Interface

In this example, the network administrator defines hierarchical queues and scheduler parameters by configuring a traffic control profile and binding it directly to a PPPoE subscriber interface. The network administrator then configures the traffic control profile on the underlying interface where a group of PPPoE interfaces reside.

This configuration is supported on the IQ2E PIC.

To use this configuration in a broadband access network, each forwarding class represents one type of services provided to a household customer and is mapped to a queue. Each PPPoE interface represents a household and provides shaping of all household traffic to an aggregate rate. The underlying logical interface where a group of PPPoE interfaces resides represents a DSLAM and provides shaping to the DSLAM rate.

Table 96 on page 1048 lists the scheduler and queue mapping for this configuration.

Table 96: Scheduler per Underlying Interface Mapping

Level	Type	Mapping
4	Queue	PPPoE interface
3	Scheduler	PPPoE interface
2	Scheduler	Underlying logical interface
1	Scheduler	Underlying interface

```
interfaces {
  ge-3/0/3 {
    hierarchical-scheduler;
    vlan-tagging;
    unit 0 {
      encapsulation ppp-over-ether;
      vlan-id 100;
    }
    unit 1 {
      vlan-id 101;
    }
  }
  pp0 {
    hierarchical-scheduler;
    unit 0 {
      pppoe-options {
        underlying-interface ge-3/0/3.0;
        server;
      }
      family inet {
        address 120.20.20.20/32 {
          destination 120.20.20.21;
        }
      }
    }
    unit 1 {
      pppoe-options {
        underlying-interface ge-3/0/3.0;
        server;
      }
      family inet {
        address 130.30.30.30/32 {
          destination 130.30.30.31;
        }
```


Copyright © 2013, Juniper Networks, Inc.

```
            }
          }
          unit 2 {
            pppoe-options {
              underlying-interface ge-3/0/3.0;
              server;
            }
            family inet {
              address 140.40.40.40/32 {
                destination 140.40.40.41;
              }
            }
          }
        }
      }
        class-of-service {
          traffic-control-profiles {
            tcp1 {
              scheduler-map data_smap;
              shaping-rate 50k;
              guaranteed-rate 10k;
            }
            tcp2 {
              scheduler-map data_smap;
              shaping-rate 50m;
              guaranteed-rate 10m;
            }
          }
          interfaces {
            pp0 {
              unit 0 {
                output-traffic-control-profile tcp1;
              }
                unit 1 {
                  output-traffic-control-profile tcp1;
                  }
                  unit 2 {
                    output-traffic-control-profile tcp1;
                    }
                    }
                    ge-3/0/3 {
                      unit 0 {
                        output-traffic-control-profile tcp2;
                      }
                    }
                  }
                  ...
                }
```

Related Documentation

- CoS for PPPoE Subscriber Interfaces Overview on page 1014
- Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface on page 1024
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring Hierarchical Scheduling and Queuing for an Interface Set of Static PPPoE Subscriber Interfaces

In this example, the network administrator defines hierarchical queues and scheduler parameters by configuring traffic-control profile and binding it directly to a PPPoE subscriber interface. The network administrator then configures the traffic-control profile on a set of PPPoE interfaces.

This configuration is supported on the IQ2E PIC.

To use this configuration in a broadband access network, each forwarding class represents one type of services provided to a household customer and is mapped to a queue. Each PPPoE interface represents a household and provides shaping of all household traffic to an aggregate rate. In addition, the PPPoE interface-set configuration provides shaping of traffic for a group of PPPoE interface on a DSLAM to a DSLAM aggregate rate.

Table 97 on page 1050 lists the scheduler and queue mapping for this configuration.

Table 97: Scheduler per Logical Interface with Interface Set Mapping

Level	Type	Mapping
4	Queue	PPPoE interface
3	Scheduler	PPPoE interface
2	Scheduler	Set of PPPoE interfaces
1	Scheduler	Underlying physical interface

```
interfaces {
  interface-set iflset1 {
    interface pp0 {
      unit 0;
      unit 1;
      unit 2;
      }
    }
    pp0 {
      unit 0 {
        pppoe-options {
          underlying-interface ge-3/0/3.0;
          server;
        }
        family inet {
          address 120.20.20.20/32 {
            destination 120.20.20.21;
          }
        }
      }
      unit 1 {
        pppoe-options {
```

 Copyright © 2013, Juniper Networks, Inc.

```
            underlying-interface ge-3/0/3.0;
            server;
          }
          family inet {
            address 130.30.30.30/32 {
              destination 130.30.30.31;
            }
          }
        }
        unit 2 {
          pppoe-options {
            underlying-interface ge-3/0/3.0;
            server;
          }
          family inet {
            address 140.40.40.40/32 {
              destination 140.40.40.41;
            }
          }
        }
      }
      ge-3/0/3 {
        hierarchical-scheduler;
        vlan-tagging;
        unit 0 {
          encapsulation ppp-over-ether;
          vlan-id 100;
        }
        unit 1 {
          vlan-id 101;
        }
        unit 2 {
          vlan-id 102;
        }
      }
    }
    class-of-service {
      traffic-control-profiles {
        tcp1 {
          scheduler-map data_smap;
          shaping-rate 50k;
          guaranteed-rate 10k;
        }
        tcp2 {
          scheduler-map data_smap;
          shaping-rate 50m;
          guaranteed-rate 10m;
        }
      }
      interfaces {
        pp0 {
          unit 0 {
            output-traffic-control-profile tcp1;
          }
          unit 1 {
            output-traffic-control-profile tcp1;
```

```
        }
        unit 2 {
            output-traffic-control-profile tcp1;
        }
        interface-set iflset1 {
            output-traffic-control-profile tcp2;
        }
        ...
    }
```

Related Documentation

- CoS for PPPoE Subscriber Interfaces Overview on page 1014
- Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface on page 1024
- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956

Example: Configuring a Dynamic Interface Set of VLAN Subscribers

- Requirements on page 1052
- Overview on page 1052
- Configuring the Dynamic VLANs on page 1052
- Configuring Dynamic Traffic Scheduling and Shaping on page 1054
- Configuring the Interface Set in the Dynamic Profile on page 1057
- Configuring DHCP Access on page 1059
- Configuring RADIUS Authentication on page 1060
- Verification on page 1064

Requirements

This example uses the following software and hardware components:

- Junos OS Release 10.4
- MX Series Router with MPCs

Overview

In this example, the network administrator groups dynamic VLAN interfaces in an interface set. The interface set is configured in a dynamic profile, and enables hierarchical scheduling for the VLAN interfaces for a multiplay service.

DHCP is used as the access method, and RADIUS is used as the authentication method for the interfaces associated with the interface set.

Configuring the Dynamic VLANs

CLI Quick Configuration

To quickly configure the dynamic VLANs, copy the following commands and paste them into the router terminal window:

```
[edit]
```


Copyright © 2013, Juniper Networks, Inc.

```
edit dynamic-profiles vlan-prof
edit interfaces $junos-interface-ifd-name unit $junos-interface-unit
set vlan-id $junos-vlan-id
set demux-source inet
set family inet unnumbered-address lo0.0 preferred-source-address 100.20.32.2
top
edit interfaces ge-1/0/0
set hierarchical-scheduler
set vlan-tagging
edit auto-configure vlan-ranges dynamic-profile vlan-prof
set ranges any
set accept inet
top
set interfaces lo0 unit 0 family inet address 100.20.32.2/32
```

Configuring the Dynamic Profile for the Autoconfigured VLANs

Step-by-Step Procedure

In this section, you create a dynamic profile for the VLAN IDs to be automatically assigned when subscribers log in.

To configure the dynamic profile for the VLANs:

1. Configure the dynamic profile.

   ```
   [edit]
   user@host#edit dynamic-profile vlan-prof
   ```

2. Configure the interfaces.

   ```
   [edit dynamic-profiles vlan-prof]
   user@host#edit interfaces $junos-interface-ifd-name unit $junos-interface-unit
   ```

3. Add the VLAN ID variable.

   ```
   [edit dynamic-profiles vlan-prof interfaces $junos-interface-ifd-name unit
     $junos-interface-unit]
   user@host#set vlan-id $junos-vlan-id
   ```

4. Configure the demux source as IPv4.

   ```
   [edit dynamic-profiles vlan-prof interfaces $junos-interface-ifd-name unit
     $junos-interface-unit]
   user@host#set demux-source inet
   ```

5. Configure the family.

   ```
   [edit dynamic-profiles vlan-prof interfaces $junos-interface-ifd-name unit
     $junos-interface-unit]
   user@host#set family inet unnumbered-address lo0.0 preferred-source-address
     100.20.32.2
   ```

Configuring the VLAN Interfaces

Step-by-Step Procedure

To configure the VLAN interfaces:

1. Create the VLAN interface.

   ```
   [edit]
   user@host# edit interfaces ge-1/0/0
   ```

2. Enable hierarchical scheduling.

```
[edit interfaces ge-1/0/0]
user@host# set hierarchical-scheduler
```

3. Configure VLAN tagging.

```
[edit interfaces ge-1/0/0]
user@host# set vlan-tagging
```

4. Configure auto-configuration for the dynamic profile.

```
[edit interfaces ge-1/0/0]
user@host# edit auto-configure vlan-ranges dynamic-profile vlan-prof
```

5. Configure any VLAN ID range.

```
[edit interfaces ge-1/0/0 auto-configure vlan-ranges dynamic-profile vlan-prof]
user@host# set ranges any
```

6. Specify IPv4 traffic for the VLAN.

```
[edit interfaces ge-1/0/0 auto-configure vlan-ranges dynamic-profile vlan-prof]
user@host# set accept inet
```

Configuring the Loopback Interface

Step-by-Step Procedure

To configure the loopback interface:

1. Create the loopback interface.

```
[edit]
user@host# edit interfaces lo0
```

2. Configure the unit and the family.

```
[edit intefaces lo0]
user@host# set unit 0 family inet address 100.20.32.2/32
```

Configuring Dynamic Traffic Scheduling and Shaping

CLI Quick Configuration

To quickly configure the traffic scheduling and shaping parameters, copy the following commands and paste them into the router terminal window:

```
[edit]
edit dynamic-profiles multiplay class-of-service schedulers be_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up
edit ef_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up
edit af_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
```


Copyright © 2013, Juniper Networks, Inc.

```
up
edit nc_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up
edit voice_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up
edit video_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up
edit game_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up
edit data_sch
set transmit-rate percent 12
set buffer-size percent 12
set priority low
up 2
edit scheduler-maps all_smap
set forwarding-class be scheduler be_sch
set forwarding-class ef scheduler ef_sch
set forwarding-class af scheduler af_sch
set forwarding-class nc scheduler nc_sch
set forwarding-class voice scheduler voice_sch
set forwarding-class video scheduler video_sch
set forwarding-class game scheduler game_sch
set forwarding-class data scheduler data_sch
up 2
edit traffic-control-profiles multiplay
set scheduler-map all_smap
set shaping-rate 100m
set guaranteed-rate 20m
```

Configuring the Schedulers in the Dynamic Profile

Step-by-Step Procedure

In this section, you create a dynamic profile for the multiplay service and configure scheduling and shaping.

To configure the schedulers:

1. Create the **multiplay** dynamic profile.

   ```
   [edit]
   user@host# edit dynamic-profiles multiplay class-of-service schedulers
   ```

2. Configure the best effort scheduler.

   ```
   [edit dynamic-profiles multiplay class-of-service schedulers]
   ```

```
user@host# edit be_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```

3. Configure the expedited forwarding scheduler.

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit ef_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```

4. Configure the assured forwarding scheduler.

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit af_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```

5. Configure the network control scheduler.

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit nc_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```

6. Configure the voice scheduler.

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit voice_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```

7. Configure the video scheduler.

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit video_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```

8. Configure the gaming scheduler.

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit game_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```

9. Configure the data scheduler.

```
[edit dynamic-profiles multiplay class-of-service schedulers]
user@host# edit data_sch
user@host# set transmit-rate percent 12
user@host# set buffer-size percent 12
user@host# set priority low
```


Copyright © 2013, Juniper Networks, Inc.

Configuring the Scheduler Map in the Dynamic Profile

Step-by-Step Procedure

To configure the scheduler map:

1. Configure the scheduler map for all of the services.

```
[edit dynamic-profiles multiplay class-of-service]
user@host# edit scheduler-maps all_smap
```

2. Configure the forwarding classes for each service in the scheduler map.

```
[edit dynamic-profiles multiplay class-of-service scheduler-maps all_smap]
user@host# set forwarding-class be scheduler be_sch
user@host# set forwarding-class ef scheduler ef_sch
user@host# set forwarding-class af scheduler af_sch
user@host# set forwarding-class nc scheduler nc_sch
user@host# set forwarding-class voice scheduler voice_sch
user@host# set forwarding-class video scheduler video_sch
user@host# set forwarding-class game scheduler game_sch
user@host# set forwarding-class data scheduler data_sch
```

Configuring the Traffic-Control Profile in the Dynamic Profile

Step-by-Step Procedure

To configure the traffic-control profile the interface set:

1. Configure the traffic-control profile.

```
[edit dynamic-profiles multiplay class-of-service]
user@host# edit traffic control-profiles multiplay
```

2. Configure the scheduler map.

```
[edit dynamic-profiles multiplay class-of-service traffic control-profiles multiplay]
user@host# set scheduler-map all_smap
```

3. Configure the shaping rate.

```
[edit dynamic-profiles multiplay class-of-service traffic control-profiles multiplay]
user@host# set shaping-rate 100m
```

4. Configure the guaranteed rate.

```
[edit dynamic-profiles multiplay class-of-service traffic control-profiles multiplay]
user@host# set guaranteed-rate 20m
```

Configuring the Interface Set in the Dynamic Profile

CLI Quick Configuration

To quickly configure the interface set, copy the following commands and paste them into the router terminal window:

```
[edit]
edit dynamic-profiles multiplay
edit interfaces interface-set $junos-interface-set-name
set interface $junos-interface-ifd-name unit $junos-underlying-interface-unit
top
edit class-of-service interfaces interface-set
set output-traffic-control-profile multiplay
```

Configuring the Interfaces for the Interface Set

Step-by-Step Procedure

To configure the interface variable for the interface set:

1. Configure the dynamic profile for the interface set.

   [edit]
   user@host#**edit dynamic-profiles multiplay**

2. Configure the interface using the Junos OS predefined variable.

   [edit dynamic-profiles multiplay]
   user@host#**edit interfaces $junos-interface-ifd-name unit $junos-underlying-interface-unit**

3. Configure the family.

   [edit dynamic-profiles multiplay interfaces $junos-interface-set-name unit $junos-underlying-interface-unit]
   user@host#**set family inet unnumbered-address lo0.0 preferred-source-address 100.20.32.2**

Configuring the Interface Set

Step-by-Step Procedure

To configure the interface set:

1. Configure the interface set using the Junos OS predefined variable.

   [edit dynamic-profiles multiplay]
   user@host#**edit interfaces interface-set $junos-interface-set-name**

2. Add the dynamic VLAN interfaces to the interface set.

   [edit dynamic-profiles multiplay interfaces $junos-interface-set-name]
   user@host#**set interface $junos-interface-ifd-name unit $junos-underlying-interface-unit**

Applying the Traffic Control Profile to the Interface Set

Step-by-Step Procedure

You apply the traffic control profile outside of the dynamic profile in the **[edit class-of-service]** hierarchy.

To apply the traffic-control profile:

1. Specify the interface set to which you want to apply the traffic control profile.

   [edit class-of-service]
   user@host#**edit interfaces interface-set dynamic-set**

2. Attach the output-traffic control profile defined in the dynamic profile to the interface set.

   [edit class-of-service interfaces]
   user@host#**set output-traffic-control-profile multiplay**

Copyright © 2013, Juniper Networks, Inc.

Configuring DHCP Access

CLI Quick Configuration

To quickly configure DHCP access, copy the following commands and paste them into the router terminal window:

```
[edit]
edit system services dhcp-local-server authentication
set password multiplay
set username-include user-prefix multiplay
up 1
set dynamic-profile dhcp-vlan-prof aggregate-clients replace
set group vlans interface ge-1/0/0
top
edit access address-assignment pool v4 family inet
set network 100.20.0.0/16
set range limited low 100.20.0.10
set range limited high 100.20.128.250
set dhcp-attributes maximum-lease-time 84600
```

Configuring the DHCP Local Server

Step-by-Step Procedure

To configure DHCP access:

1. Configure the DHCP local server.

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

2. Set the password.

   ```
   [edit system services dhcp-local-server authentication]
   user@host# set password multiplay
   ```

3. Specify that you want to include optional information in the username.

   ```
   [edit system services dhcp-local-server authentication]
   user@host# set username-include user-prefix multiplay
   ```

4. Attach the dynamic profile with the interface set.

   ```
   [edit system services dhcp-local-server]
   user@host# set dynamic-profile dhcp-vlan-prof aggregate-clients replace
   ```

5. Configure a group for the VLAN interface.

   ```
   [edit system services dhcp-local-server]
   user@host# set group vlans interface ge-1/0/0
   ```

Configuring Address Assignment Pools

Step-by-Step Procedure

To configure address assignment pools:

1. Configure the pool of IPv4 addresses.

   ```
   [edit access]
   user@host#edit address-assignment pool v4 family inet
   ```

2. Configure the family of interfaces in the pool.

```
[edit access address-assignment pool v4]
user@host#set network 100.20.0.0/16
```

3. Configure the upper and lower bounds of the address range.

```
[edit access address-assignment pool v4]
user@host#set range limited low 100.20.0.10
user@host#set range limited high 100.20.128.250
```

4. Configure the maximum length of time in seconds for which a subscriber can request and hold a lease.

```
[edit access address-assignment pool v4]
user@host#set dhcp-attributes maximum-lease-time 84600
```

Configuring RADIUS Authentication

CLI Quick Configuration

To quickly configure RADIUS authentication, copy the following commands and paste them into the router terminal window:

```
[edit]
edit access radius-server 172.28.30.108
set secret $9$1u5ErvW87bwgSr4Zji5T
set timeout 5
set retry 5
up 2
edit profile acc-prof
set authentication-order radius
set radius authentication-server 172.28.30.108
```

Configuring RADIUS Access

Step-by-Step Procedure

To configure RADIUS access:

1. Configure the RADIUS server.

```
[edit access]
user@host#edit radius-server 172.28.30.108
```

2. Configure the required secret (password) that the local router or switch passes to the RADIUS client.

```
[edit access radius-server 172.28.30.108]
user@host# set secret $9$1u5ErvW87bwgSr4Zji5T
```

3. Configure the length of time that the local router or switch waits to receive a response from a RADIUS server.

```
[edit access radius-server 172.28.30.108]
user@host# set timeout 5
```

4. Configure the number of times that the router or switch attempts to contact a RADIUS accounting server.

```
[edit access radius-server 172.28.30.108]
user@host# set retry 5
```

5. Configure the access profile.

Copyright © 2013, Juniper Networks, Inc.

```
[edit access]
user@host#edit profile acc-prof
```

6. Configure the authentication order.

```
[edit access profile acc-prof ]
user@host# set authentication-order radius
```

7. Configure the authentication server.

```
[edit access profile acc-prof]
user@host#set radius authentication-server 172.28.30.108
```

Results

```
dynamic-profiles {
  vlan-prof {
    interfaces {
      “$junos-interface-ifd-name” {
        unit "$junos-interface-unit" {
          vlan-id "$junos-vlan-id";
          demux-source inet;
          family inet {
            unnumbered-address lo0.0 preferred-source-address 100.20.32.2;
          }
        }
      }
    }
  }
  multiplay {
    class-of-service {
      traffic-control-profiles {
        multiplay {
          scheduler-map all_smap;
          shaping-rate 100m;
          guaranteed-rate 20m;
        }
      }
      interfaces {
          interface-set “$junos-interface-set-name” {
            interface “$junos-interface-ifd-name” {
              unit “$junos-underlying-interface-unit”;
            }
          }
        “$junos-interface-ifd-name” {
          unit "$junos-interface-unit" {
            output-traffic-control-profile multiplay;
          }
        }
      }
      scheduler-maps {
        all_smap {
          forwarding-class be scheduler be_sch;
          forwarding-class ef scheduler ef_sch;
          forwarding-class af scheduler af_sch;
          forwarding-class nc scheduler nc_sch;
          forwarding-class voice scheduler voice_sch;
```

```
            forwarding-class video scheduler video_sch;
            forwarding-class game scheduler game_sch;
            forwarding-class data scheduler data_sch;
        }
    }
    schedulers {
        be_sch {
            transmit-rate percent 12;
            buffer-size percent 12;
            priority low;
        }
        ef_sch {
            transmit-rate percent 12;
            buffer-size percent 12;
            priority low;
        }
        af_sch {
            transmit-rate percent 12;
            buffer-size percent 12;
            priority low;
        }
        nc_sch {
            transmit-rate percent 12;
            buffer-size percent 12;
            priority low;
        }
        voice_sch {
            transmit-rate percent 12;
            buffer-size percent 12;
            priority low;
        }
        video_sch {
            transmit-rate percent 12;
            buffer-size percent 12;
            priority low;
        }
        game_sch {
            transmit-rate percent 12;
            buffer-size percent 12;
            priority low;
        }
        data_sch {
            transmit-rate percent 12;
            buffer-size percent 12;
            priority low;
        }
    }
    }
}
access {
    radius-server {
        172.28.30.108 {
            secret "$9$1u5ErvW87bwgSr4Zji5T"; ## SECRET-DATA
            timeout 5;
            retry 5;
        }
```


Copyright © 2013, Juniper Networks, Inc.

```
            }
            profile acc-prof {
                authentication-order radius;
                radius {
                    authentication-server 172.28.30.108;
                }
            }
            address-assignment {
                pool v4 {
                    family inet {
                        network 100.20.0.0/16;
                        range limited {
                            low 100.20.0.10;
                            high 100.20.128.250;
                        }
                        dhcp-attributes {
                            maximum-lease-time 84600;
                        }
                    }
                }
            }
        }
        class-of-service {
            interfaces {
                interface-set dynamic-set {
                    output-traffic-control-profile multiplay;
                }
            }
        }
        interfaces {
            interface-set “$junos-interface-set-name” {
                interface "$junos-interface-ifd-name" {
                    unit "$junos-underlying-interface-unit";
                }
            }
            "$junos-interface-ifd-name" {
                unit "$junos-underlying-interface-unit" {
                    family inet {
                        unnumbered-address lo0.0 preferred-source-address 100.20.32.2;
                    }
                }
            }
        }
    }
}
    interfaces {
        ge-1/0/0 {
            hierarchical-scheduler;
            vlan-tagging;
            auto-configure {
                vlan-ranges {
                    dynamic-profile vlan-prof {
                        accept inet;
                        ranges {
                            any;
                        }
```

```
            }
          }
        }
      }
      lo0 {
        unit 0 {
          family inet {
            address 100.20.32.2/32;
          }
        }
      }
    }
    system {
      services {
        dhcp-local-server {
          authentication {
            password multiplay;
            username-include {
              user-prefix multiplay;
            }
          }
          dynamic-profile multiplay aggregate-clients replace;
          group vlans {
            interface ge-1/0/0.0;
          }
        }
      }
    }
```

Verification

To confirm that the configuration is correct, perform these tasks:

- Verifying the Interfaces that are Included in the Interface Set on page 1064
- Verifying the Traffic Scheduling and Shaping Parameters for the Interface Set on page 1064

Verifying the Interfaces that are Included in the Interface Set

Purpose Verify the interfaces included in the interface set.

Action `user@host>` **show interfaces interface-set dynamic-set terse**

Verifying the Traffic Scheduling and Shaping Parameters for the Interface Set

Purpose Verify that the traffic scheduling and shaping parameters are applied properly to an interface included in the interface set.

Action `user@host>` **show class-of-service interface**

Related Documentation

- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Configuring an Interface Set of Subscribers in a Dynamic Profile on page 1028

 Copyright © 2013, Juniper Networks, Inc.

Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile

Interface sets enable you to provide hierarchical scheduling to a group of subscriber interfaces. In this example, by using the **$junos-svlan-interface-set-name** internal dynamic variable when specifying the interface set name, you can locally generate an interface set name for use by SVLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is ***physical_interface_name - outer_VLAN_tag***.

- Requirements on page 1065
- Overview on page 1065
- Configuration on page 1066
- Verification on page 1068

Requirements

Dynamic SVLAN traffic shaping is supported only on Juniper Networks MX Series 3D Universal Edge Routers running Junos OS Release 11.4R2 or later.

Before you begin, configure the subscriber interfaces that you intend to include in the interface set. You can find general configuration instructions for the supported dynamic interface configuration in the *Junos OS Subscriber Management and Services Library* as follows:

- For dynamic VLAN interfaces, see "Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet" on page 824.
- For dynamic IP demux interfaces, see "Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles" on page 771 and "Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet" on page 825.
- For dynamic VLAN demux interfaces, see "Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles" on page 772.

Overview

Interface sets enable you to provide hierarchical scheduling to a group of subscriber interfaces. By using the **$junos-svlan-interface-set-name** internal dynamic variable when specifying the interface set name, you can locally generate an interface set name for use by SVLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is ***physical_interface_name - outer_VLAN_tag***.

This example includes the following statements:

- **interface-set**—Configures the name of the scheduler for dynamic CoS. In this example, you use the **$junos-svlan-interface-set-name** variable to obtain the locally generated interface set name for use by SVLAN interfaces based on the outer tag of the dual-tagged VLAN.
- **output-traffic-control-profile**—Applies an output traffic scheduling and shaping profile to the interface set.

Copyright © 2013, Juniper Networks, Inc.

- **output-traffic-control-profile-remaining**—Applies an output traffic scheduling and shaping profile for remaining traffic to the interface set.

Configuration

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]
set dynamic-profiles profile-dhcp-ipdemux interfaces interface-set
   $junos-svlan-interface-set-name interface $junos-interface-ifd-name unit
   $junos-underlying-interface-unit
set dynamic-profiles profile-dhcp-ipdemux interfaces $junos-interface-ifd-name unit
   $junos-underlying-interface-unit
set class-of-service traffic-control-profiles tcp1 scheduler-map schedMap
set class-of-service traffic-control-profiles tcp1 shaping-rate 50m
set class-of-service traffic-control-profiles tcp1 guaranteed-rate 200k
set class-of-service traffic-control-profiles tcp3 scheduler-map ss1q0q1
set class-of-service traffic-control-profiles tcp3 shaping-rate 20m
set class-of-service traffic-control-profiles tcp3 guaranteed-rate 5m
set class-of-service interfaces interface-set ae0-111 output-traffic-control-profile tcp1
set class-of-service interfaces interface-set ae0-111
   output-traffic-control-profile-remaining tcp3
```

Step-by-Step Procedure

To configure an SVLAN interface set of subscriber interfaces:

1. Access the dynamic profile you want to modify for interface sets.

   ```
   [edit]
   user@host# edit dynamic-profiles profile-dhcp-ipdemux
   ```

2. Access the dynamic profile interface configuration.

   ```
   [edit dynamic-profiles profile-dhcp-ipdemux]
   user@host# edit interfaces
   ```

3. Configure the SVLAN interface set in the dynamic profile.

   The interface set is created dynamically when the subscriber logs in.

   ```
   [edit dynamic-profiles profile-dhcp-ipdemux interfaces]
   user@host# edit interface-set $junos–svlan-interface-set-name
   ```

4. Include dynamic IP demux interface creation within the dynamic interface set.

   ```
   [edit dynamic-profiles profile-dhcp-ipdemux interfaces interface-set
      $junos-svlan-interface-set-name]
   user@host# set interface $junos-interface-ifd-name unit
      $junos-underlying-interface-unit
   ```

5. Access the SVLAN interface set name that you expect **$junos-svlan-interface-set-name** to generate. For example, to specify the expected interface set name for aggregated Ethernet interface ae0 and outer VLAN tag 111, include **ae0-111** for the ***interface-set-name*** variable.

   ```
   [edit class-of-service interfaces]
   ```


Copyright © 2013, Juniper Networks, Inc.

```
user@host# edit interface-set ae0-111
```

6. Apply traffic shaping and queuing parameters to the SVLAN interface set.

> TIP: You must configure the interface set in the static **[edit class-of-service]** hierarchy, not in the **[edit dynamic-profiles]** hierarchy.

```
[edit class-of-service interfaces interface-set ae0-111]
user@host# set output-traffic-control-profile tcp1
```

7. Apply traffic shaping and queuing parameters to any remaining traffic on the SVLAN interface set.

```
[edit class-of-service interfaces interface-set ae0-111]
user@host# set output-traffic-control-profile-remaining tcp3
```

Results

From configuration mode, confirm your configuration by entering the **show dynamic-profiles** command and the **show class-of-service** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show dynamic-profiles
dynamic-profiles {
   profile-dhcp-ipdemux {
      interfaces {
         interface-set "$junos-svlan-interface-set-name" {
            interface "$junos-interface-ifd-name" {
               unit "$junos-underlying-interface-unit";
            }
         }
         "$junos-interface-ifd-name" {
            unit "$junos-underlying-interface-unit";
         }
      }
   }
}
```

```
user@host# show class-of-service
class-of-service {
   traffic-control-profiles {
      tcp1 {
         scheduler-map schedMap;
         shaping-rate 50m;
         guaranteed-rate 200k;
      }
      tcp3 {
         inactive: scheduler-map ss1q0q1;
         shaping-rate 20m;
         guaranteed-rate 5m;
      }
   }
```

```
    interfaces {
      interface-set ae0-111 {
        output-traffic-control-profile tcp1;
        output-traffic-control-profile-remaining tcp3;
      }
    }
  }
```

Verification

To confirm that the configuration is correct, perform these tasks:

Verifying the Interfaces that are Included in the Interface Set

Purpose Verify the interfaces that are included in the interface set.

Action

```
user@host> show class-of-service interface-set
```

Displaying Information for Active Subscribers

Purpose Display information for active subscribers.

Action

```
user@host> show subscribers detail
```

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring a Basic Dynamic Profile on page 673
- *Configuring Hierarchical Schedulers for CoS*
- Configuring Remaining Common Queues on MIC and MPC Interfaces on page 1094

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 59

Bandwidth Management for Dynamic CoS Overview

- Excess Bandwidth Distribution on MIC and MPC Interfaces Overview on page 1069
- Traffic Burst Management on MIC and MPC Interfaces Overview on page 1070
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072
- Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 1074
- Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview on page 1076
- Hierarchical CoS Shaping-Rate Adjustments Overview on page 1080
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 1082
- Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 1083
- CoS Adjustment Control Profiles Overview on page 1084

Excess Bandwidth Distribution on MIC and MPC Interfaces Overview

Service providers often used tiered services to provide bandwidth for excess traffic as traffic patterns vary. By default, excess bandwidth between a configured guaranteed rate and shaping rate is shared equally among all queues on MIC and MPC interfaces, which might not be optimal for all subscribers to a service.

You can adjust this distribution by configuring the rates and priorities for the excess bandwidth.

By default, when traffic exceeds the shaping or guaranteed rates, the system demotes traffic with guaranteed high (GH) priority and guaranteed medium (GM) priority. You can disable this priority demotion for the MIC and MPC interfaces in your router.

Related Documentation

- *Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs*
- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 1090
- *Per-Priority Shaping on MIC and MPC Interfaces Overview*
- Traffic Burst Management on MIC and MPC Interfaces Overview on page 1070

Copyright © 2013, Juniper Networks, Inc.

Traffic Burst Management on MIC and MPC Interfaces Overview

You can manage the impact of bursts of traffic on your network by configuring a burst-size value with the shaping rate or the guaranteed rate. The value is the maximum bytes of rate credit that can accrue for an idle queue or scheduler node. When a queue or node becomes active, the accrued rate credits enable the queue or node to catch up to the configured rate.

Figure 30: Sample Burst Shaping Rates

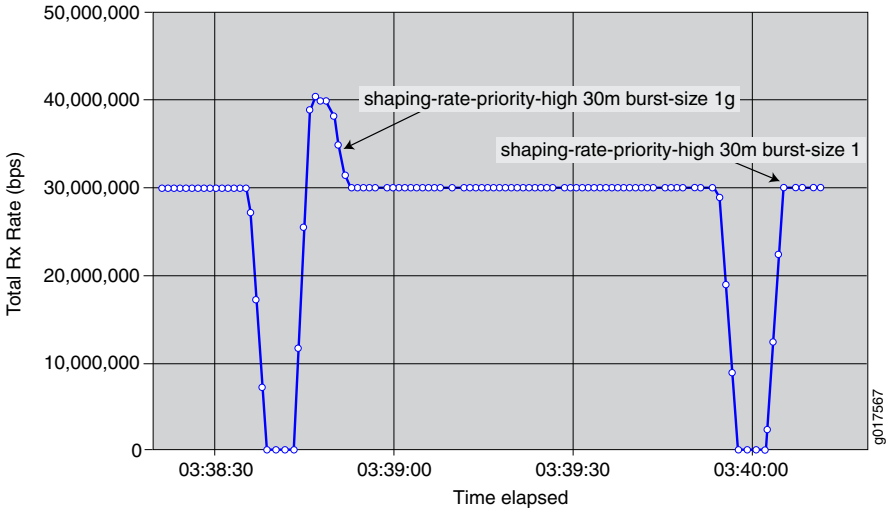

In Figure 30 on page 1070, the network administrator configures a large burst-size value for the shaping rate, then configures a small burst-size value. The larger burst size is subject to a maximum value. The smaller burst size is subject to a minimum value that enables the system to achieve the configured rates.

In both configurations, the scheduler node can burst beyond its shaping rate for a brief interval. The burst of traffic beyond the shaping rate is more noticeable with the larger burst size than the smaller burst size.

- Guidelines for Configuring the Burst Size on page 1070
- How the System Calculates the Burst Size on page 1071

Guidelines for Configuring the Burst Size

Typically, the default burst-size (100 ms) for both scheduler nodes and queues on MIC and MPC interfaces is adequate for most networks. However, if you have intermediate equipment in your network that has very limited buffering and is intolerant of bursts of traffic, you might want to configure a lower value for the burst size.

Use caution when selecting a different burst size for your network. A burst size that is too high can overwhelm downstream networking equipment, causing dropped packets and inefficient network operation. Similarly, a burst size that is too low can prevent the network from achieving your configured rate.

Copyright © 2013, Juniper Networks, Inc.

When configuring a burst size, keep the following considerations in mind:

- The system uses an algorithm to determine the actual burst size that is implemented for a node or queue. For example, to reach a shaping rate of 8 Mbps, you must allocate 1Mb of rate credits every second. A shaping rate of 8 Mbps with a burst size of 500,000 bytes of rate-credit per seconds enables the system to transmit at most 500,000 bytes, or 4 Mbps. The system cannot implement a burst size that prevents the rate from being achieved.

  For more information, see "How the System Calculates the Burst Size" on page 1071.

- There are minimum and maximum burst sizes for each platform, and different nodes and queue types have different scaling factors. For example, the system ensures the burst cannot be set lower than 1 Mbps for a shaping rate of 8 Mbps. To smoothly shape traffic, rate credits are sent much faster than once per second. The interval at which rate credits are sent varies depending on the platform, the type of rate, and the scheduler level.
- When you have configured adjustments for the shaping rate (either by percentage or through an application such as ANCP or Multicast OIF), the system bases the default and minimum burst-size calculations on the adjusted shaping rate.
- When you have configured cell shaping mode to account for ATM cell tax, the system bases the default and minimum burst-size calculations on the post-tax shaping rate.
- The guaranteed rate and shaping rate share the value specified for the burst size. If the guaranteed rate has a burst size specified, that burst size is used for the shaping rate; if the shaping rate has a burst size specified, that bursts size is used for the guaranteed rate. If you have specified a burst size for both rates, the system uses the lesser of the two values.
- The burst size configured for the guaranteed rate cannot exceed the burst-size configured for the shaping rate. The system generates a commit error.
- If you have not configured a guaranteed rate, logical interfaces and interface sets receive a default guaranteed rate from the port speed. Queues receive a default guaranteed rate from the parent logical interface or interface set.

How the System Calculates the Burst Size

When calculating the burst size, the system uses an exponent of a power of two. For example:

Shaping-rate in bps * 100 ms / (8 bits/byte * 1000 ms/s) = 1,875,000 bytes

The system then rounds this value up. For example, the system uses the following calculation to determine the burst size for a scheduler node with a shaping rate of 150 Mbps:

Max (Shaping rate, Guaranteed rate) bps * 100 ms / (8 bits/byte * 1000 ms/s) = 1,875,000 bytes

Rounded up to the next higher power of two = 2,097,150 (which is 221, or 0x200000)**

The system assigns a single burst size to each of the following rate pairs:

- Shaping rate and guaranteed rate
- Guaranteed high (GH) and guaranteed medium (GM)
- Excess high (EH) and excess low (EL)
- Guaranteed low (GL)

To calculate the burst size for each pair, the system:

- Uses the configured burst-size if only one of the pair is configured.
- Uses the lesser of the two burst sizes if both values are configured.
- Uses the next lower power of two.
- To calculate the minimum burst size, the system uses the greater of the two rates.

Related Documentation

- *Per-Priority Shaping on MIC and MPC Interfaces Overview*
- *Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs*

Bandwidth Management for Downstream Traffic in Edge Networks Overview

In a subscriber access network, traffic with different encapsulations can be passed downstream to other customer premise equipment (CPE) through the MX Series router. Managing the bandwidth of downstream ATM traffic to Ethernet interfaces can be especially difficult because of the different Layer 2 encapsulations.

The *overhead accounting* feature enables you to shape traffic based on either frames or cells and assign a byte adjustment value to account for different encapsulations.

This feature is available on MIC and MPC interfaces.

Guidelines for Configuring the Shaping Mode

Frame shaping mode is useful for adjusting downstream traffic with different encapsulations. In frame shaping mode, shaping is based on the number of bytes in the frame, without regard to cell encapsulation or padding overhead. Frame is the default shaping mode on the router.

Cell shaping mode is useful for adjusting downstream cell-based traffic. In cell shaping mode, shaping is based on the number of bytes in cells, and accounts for the cell encapsulation and padding overhead.

When you specify cell mode, the resulting traffic stream conforms to the policing rates configured in downstream ATM switches, reducing the number of packet drops in the Ethernet network.

To account for ATM segmentation, the MX Series router adjusts all of the rates by 48/53 to account for ATM AAL5 encapsulation. In addition, the router accounts for cell padding, and internally adjusts each frame by 8 bytes to account for the ATM trailer.

 Copyright © 2013, Juniper Networks, Inc.

Guidelines for Configuring Byte Adjustments

When the downstream traffic has different byte sizes per encapsulation, it is useful to configure a *byte adjustment* value to adjust the frame sizes. For example, you can configure the frame shaping mode and a byte adjustment value to account for differences in Layer 2 protocols for downstream Ethernet traffic.

We recommend that you specify a byte adjustment value that represents the difference between the CPE protocol overhead and B-RAS protocol overhead.

The system rounds up the byte adjustment value to the nearest multiple of 4. For example, a value of 6 is rounded to 8, and a value of –10 is rounded to –8.

You do not need to configure a byte adjustment value to account for the downstream ATM network. However, you can specify the byte value to account for additional encapsulations or decapsulations in the downstream network.

Relationship with Other CoS Features

Enabling the overhead accounting feature affects the resulting shaping rates, guaranteed rate, and excess rate parameters, if they are configured.

The overhead accounting feature also affects the egress shaping overhead feature that you can configure at the chassis level. We recommend that you use the egress shaping-overhead feature to account for the Layer 2 overhead of the outgoing interface, and use the overhead-accounting feature to account for downstream traffic with different encapsulations and cell-based networks.

When both features are configured together, the total byte adjustment value is equal to the adjusted value of the overhead-accounting feature plus the value of the egress-shaping-overhead feature. For example, if the configured byte adjustment value is 40, and the router internally adjusts the size of each frame by 8, the adjusted overhead accounting value is 48. That value is added to the egress shaping overhead of 24 for a total byte adjustment value of 72.

Setting the Shaping-Rate and Overhead-Accounting Class-of-Service Attributes Based on Access Line Information

You can also use access line parameters in PPPoE discovery packets to set the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network. This feature is supported on MIC and MPC 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).

You can mix ANCP and PPPoE vendor-specific tags for dynamically instantiated static and interface sets so that the shaping rate is first set using PPPoE vendor-specific tags and is later adjusted by ANCP. In this case, the shaping rate value overrides the PPPoE value.

Copyright © 2013, Juniper Networks, Inc.

The Access Node or DSLAM may forward access line information using several methods. This feature only uses access line information received in Vendor-Specific Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101].

When you enable this feature, the values supplied by the PPPoE vendor-specific tags override the parameters that you have configured in the CLI for shaping-rate and overhead-accounting statements at the **[edit dynamic-profiles profile-name class-of-service traffic-control-profile]** hierarchy level. The shaping rate is based on the actual-data-rate-downstream attribute, and is only overridden if the vs-tag value is less than the configured value.

Related Documentation

- To configure overhead accounting for static Ethernet interfaces, see *Configuring Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates*
- To configure overhead accounting for dynamic subscriber access, see Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 1092
- Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 1074
- Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces on page 1093

Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags

You can use access line parameters in PPPoE discovery packets to set the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network. 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).

You can configure class-of-service attributes, for example the shaping-rate, using the CLI, RADIUS vendor-specific attributes, ANCP, multicast, or in this case, PPPoE vendor-specific tags.

CLI Interaction with PPPoE Vendor-Specific Tags

When you enable this feature, the values supplied by the PPPoE vendor-specific tags override the parameters that you have configured in the CLI for the **shaping-rate** and **overhead-accounting** statements at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles]** hierarchy level. The shaping rate is based on the actual-data-rate-downstream attribute, and is only overridden if the vs-tag value is less than the configured value.

To enable this feature, include the **actual-data-rate-downstream** or **access-loop-encapsulation** option with the **vendor-specific-tags** statement at the **[edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]** hierarchy level.

 Copyright © 2013, Juniper Networks, Inc.

RADIUS Interaction with PPPoE Vendor-Specific Tags

When you enable this feature, the PPPoE vendor-specific tags override the dynamic configuration of the shaping-rate and overhead-accounting values in RADIUS vendor-specific attributes. The shaping-rate value is only overridden if the vs-tag value is less than the RADIUS value.

RADIUS COA can overwrite the existing values. Upon receipt of a RADIUS COA, the RADIUS value overrides the value set from the PPPoE vendor-specific tags.

PPPoE vendor-specific tags can override the RADIUS values, but a later RADIUS COA request can then override that value.

ANCP Interaction with PPPoE Vendor-Specific Tags

You can mix ANCP and PPPoE vendor-specific tags on dynamic PPPoE interfaces, dynamically instantiated PPPoE interfaces, and ACI-sets. ANCP values override the PPPoE values. In this case, the ANCP shaping rate value overrides the PPPoE value.

Multicast QoS Adjustment Interaction with PPPoE Vendor-Specific Tags

Multicast QoS adjustments are not affected by this feature. The multicast adjustments adjust the shaping-rate set by PPPoE vendor-specific tags.

Shaping Rate Restrictions

Shaping rate has the following restrictions regarding the downstream-rate:

- If the downstream-rate is less than the configured shaping-rate (as set in the CLI or using RADIUS attributes) then it is applied, subject to other restrictions. If the downstream-rate is greater than or equal to the configured shaping-rate, no changes are performed.
- The downstream-rate cannot be less than a configured guaranteed-rate. If it is, the downstream-rate is set to the guaranteed-rate.
- The downstream-rate cannot be less than a configured adjust-minimum-rate. If it is, the downstream-rate is set to the adjust-mimimum-rate.
- The downstream-rate cannot be less than 1000 bps. If it is, the downstream-rate is set to 1000 bps.
- The downstream-rate cannot be less than the sum of the transmit-rates of all queues.

Related Documentation

- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072
- Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces on page 1093

Copyright © 2013, Juniper Networks, Inc.

Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview

The 30-Gigabit Ethernet Queuing and 60-Gigabit Ethernet Queuing and Enhanced Queuing Ethernet Modular Port Concentrators (MPCs) provide a set of dedicated queues for subscriber interfaces configured with hierarchical scheduling or per-unit scheduling.

The dedicated queues offered on these MPCs enable service providers to reduce costs through different scaling configurations. For example, the 60-Gigabit Ethernet Enhanced Queuing MPC enables service providers to reduce the cost per subscriber by allowing many subscriber interfaces to be created with four or eight queues. Alternatively, the 30-Gigabit Ethernet and 60-Gigabit Ethernet Queuing MPCs enable service providers to reduce hardware costs, but allow fewer subscriber interfaces to be created with four or eight queues.

This topic describes the overall queue, scheduler node, and logical interface scaling for subscriber interfaces created on these MIC and MPC combinations.

Queue Scaling for MIC and MPC Combinations

Table 98 on page 1076 lists the number of dedicated queues and number of subscribers supported per MPC.

Table 98: Dedicated Queues for MIC and MPC Interfaces

MPC	Dedicated Egress Queues	Supported Subscriber Interfaces	Logical Interfaces with 4 Queues	Logical Interfaces with 8 Queues
30-Gigabit Ethernet Queuing MPC	64,000	16,000	16,000 (8000 per PIC)	8000 (4000 per PIC)
60-Gigabit Ethernet Queuing MPC	128,000	32,000	32,000 (8000 per PIC)	16,000 (4000 per PIC)
60-Gigabit Ethernet Enhanced Queuing MPC	512,000	64,000	64,000 (16,000 per PIC)	64,000 (16,000 per PIC)

MPCs vary in the number of Packet Forwarding Engines on board. MPC1s, such as the 30-Gigabit Ethernet MPC, have one Packet Forwarding Engine. MPC2s, such as the 60-Gigabit Ethernet MPC, have two Packet Forwarding Engines. Each Packet Forwarding Engine has two schedulers that share the management of the queues.

A scheduler maps to one-half of a MIC; in CLI configuration statements, that one-half of a MIC corresponds to PIC 0, 1, 2, or 3. MIC ports are partitioned equally across the PICs. A two-port MIC has one port per PIC. A four-port MIC has two ports per PIC.

Each interface-set uses eight queues from total available egress queues.

Copyright © 2013, Juniper Networks, Inc.

Distribution of Queues on 30-Gigabit Ethernet Queuing MPCs

On 30-Gigabit Ethernet Queuing MPCs, each scheduler maps to different PICs. When only one MIC is installed, scheduler 0 maps to PIC 0 and scheduler 1 maps to PIC 1 on the MIC. When two MICs are installed, scheduler 0 can additionally distribute queues to PIC 2 on MIC 1, and scheduler 1 can additionally distribute queues to PIC 3 on MIC 1. However, the distribution of queues to the MICs is not hard-partitioned for 30-Gigabit Ethernet Queuing MPCs or other MPC1s. Distribution depends instead on how you allocate the queues to the PICs.

Figure 31 on page 1077 shows the queue distribution on a 30-Gigabit Ethernet Queuing MPC with only one MIC installed. All 64,000 egress queues on the MPC are available to the single Packet Forwarding Engine. On the Packet Forwarding Engine, half of these queues (32,000) are managed by each scheduler. Scheduler 0 contributes all of its 32,000 queues to PIC 0. Scheduler 1 contributes all of its 32,000 queues to PIC 1.

Figure 31: Distribution of Queues on the 30-Gigabit Ethernet Queuing MPC with One MIC

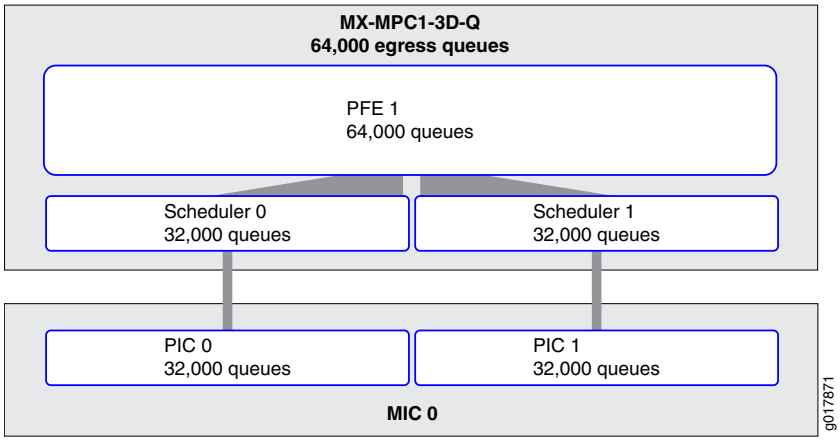

Figure 32 on page 1078 shows the queue distribution on the same MPC with two MICs installed. In this case, each scheduler can supply two PICS, one on each MIC. Because the distribution of the queues across the MICs is not hard-partitioned, you can allocate from 0 to 32,000 queues from each scheduler's pool across the scheduler's associated PICs. For example, you can allocate 32,000 queues from Scheduler 0 to PIC 0, 4000 queues from Scheduler 1 to PIC 1, and 28,000 queues from Scheduler 1 to PIC 3. Alternatively, you can allocate the queues evenly across the PICs, or allocate them in other combinations with the limitation of 32,000 queues per PIC and 32,000 queues per port.

Copyright © 2013, Juniper Networks, Inc.

Figure 32: Distribution of Queues on the 30-Gigabit Ethernet Queuing MPC with Two MICs

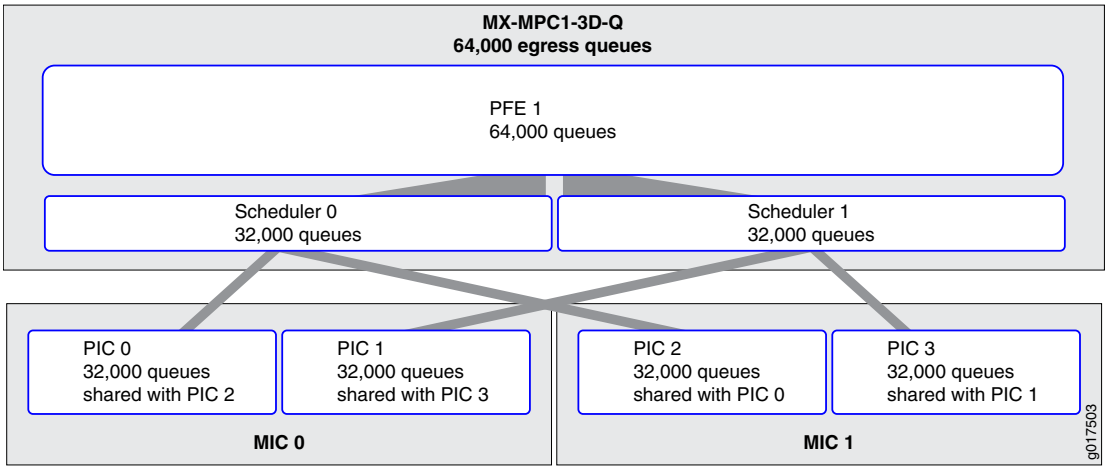

Distribution of Queues on 60-Gigabit Ethernet MPCs

On 60-Gigabit Ethernet Queuing and Enhanced Queuing Ethernet MPCs, each scheduler maps to a single PIC: PIC 0 or PIC 1 on MIC 0 and PIC 2 or PIC 3 on MIC 1. The distribution of the queues is hard-partitioned for these MPCs and other MPC2s; the only difference in distribution is in the total number of queues available.

For example, Figure 33 on page 1078 shows how queues are distributed on a 60-Gigabit Ethernet Enhanced Queuing MPC. Of the 512,000 egress queues on the MPC, half (256,000) are available to each of the two Packet Forwarding Engines. On each Packet Forwarding Engine, half of these queues (128,000) are managed by each scheduler. The complete scheduler complement (128,000) is available to only one PIC in a MIC. Thus the total number of queues available depends on the number of MICs installed. The MPC must have 2 MICs to achieve the maximum of 512,000 queues. With a single MIC, the MPC can achieve only 256,000 queues.

Figure 33: Distribution of Queues on the 60-Gigabit Ethernet Enhanced Queuing MPC

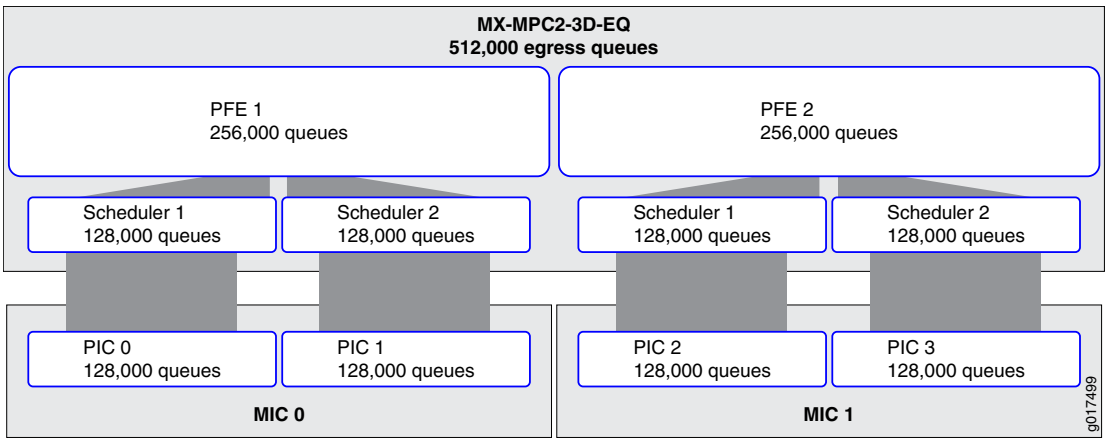

Copyright © 2013, Juniper Networks, Inc.

Determining Maximum Egress Queues and Subscriber Interfaces per Port

The number of MICs installed in an MPC and the number of ports per MIC do not affect the maximum number of queues available on a given port. These factors affect only how you are able to allocate queues (and, therefore, subscribers) for your network.

For example, a 30-Gigabit Ethernet Queuing MPC supports a maximum of 16,000 subscriber interfaces and has a maximum of 32,000 queues available per PIC. On this card, you can allocate up to 32,000 queues to a single port in each PIC. If you dedicate 4 queues per subscriber interface, you can accommodate a maximum of 8000 subscriber interfaces on a single port, and therefore need at least two ports to reach the maximum 16,000 subscriber interfaces. If you dedicate 8 queues per subscriber interface, you can accommodate a maximum of 4000 subscriber interfaces on a single port, and you need 4 ports for the maximum of 16,000 subscriber interfaces.

The 60-Gigabit Ethernet Enhanced Queuing MPC supports a maximum of 64,000 subscriber interfaces and has a maximum of 128,000 queues per PIC. You can allocate up to 128,000 queues to a single port in each PIC. However, if you dedicate 4 queues per subscriber interface, you can accommodate a maximum of only 16,000 subscriber interfaces on a single MPC port—not 32,000—because the 60-Gigabit Ethernet Enhanced Queuing MPC is limited to 16,000 subscriber interfaces per PIC. If you dedicate 8 queues per subscriber interface, you can also accommodate a maximum of 16,000 subscriber interfaces on a single MPC port. In either case, you need at least 4 ports to reach the maximum of 64,000 subscriber interfaces.

Managing Remaining Queues

When the number of available dedicated queues on the MPC drops below 10 percent, an SNMP trap is generated to notify you .

When the maximum number of dedicated queues on the MPCs is reached, a system log message, **COSD_OUT_OF_DEDICATED_QUEUES**, is generated. The system does not provide subsequent subscriber interfaces with a dedicated set of queues. For per-unit scheduling configurations, there are no configurable queues remaining on the MPC.

For hierarchical scheduling configurations, remaining queues are available when the maximum number of dedicated queues is reached on the MPC. Traffic from these logical interfaces are considered unclassified and attached to a common set of queues that are shared by all subsequent logical interfaces. These common queues are the default port queues that are created for every port. You can configure a traffic control profile and attach that to the interface to provide CoS parameters for the remaining queues.

For example, when the 30-Gigabit Ethernet Queuing MPC is configured with 32,000 subscriber interfaces with four queues per subscriber, the MPC can support 16,000 subscribers with a dedicated set of queues. You can provide CoS shaping and scheduling parameters to the remaining queues for those subscriber interfaces by attaching a special traffic-control profile to the interface.

These subscriber interfaces remain with this traffic control profile, even if dedicated queues become available.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- For information about managing dedicated queues in a static CoS configuration, see *Managing Dedicated and Remaining Queues for Static CoS Configurations on MIC and MPC Interfaces*
- For information about managing dedicated queues in a dynamic subscriber access configuration, see Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces on page 1093
- *Scheduler Node Scaling on MIC and MPC Interfaces Overview*
- *COSD System Log Messages*

Hierarchical CoS Shaping-Rate Adjustments Overview

This overview describes how MX Series 3D Universal Edge Routers installed in a subscriber access network can adjust hierarchical class-of-service (CoS) parameters to prevent bandwidth contention at subscriber interfaces.

Hierarchical CoS is supported only for subscriber interfaces on EQ DPC or MPC interfaces operating in hierarchical scheduler mode.

The characteristics of voice, data, and video applications vary widely in their requirements for traffic throughput, bandwidth management, delay and jitter tolerance, and buffer depth. To prevent bandwidth contention at subscriber interfaces, you can configure applications such as ANCP and Multicast to perform real-time adjustments to the shaping rate configured for subscriber interfaces for residential gateways. Enabling shaping-rate adjustments on the router can prevent bandwidth contention at the interface from causing degradation of the subscriber's voice, data, or video services.

Depending on the application, shaping-rate adjustments are supported on Enhanced Queueing (EQ) DPCs on MX Series routers and MPC/MIC modules on MX Series routers.

Types of Shaping-Rate Adjustments

The ANCP application supports *absolute* adjustments to a specific shaping-rate value. You can configure ANCP to communicate the subscriber local loop speed to the MX Series router, which in turn throttles traffic destined to the associated subscriber interface so that it matches the subscriber local loop speed. ANCP acquires subscriber line rate information from DSLAMs and then communicates this data transmission rate for use with CoS.

The OIF mapping and reverse OIF mapping multicast applications support *delta* adjustments that increase or decrease the current shaping rate by a certain value. The system adjusts traffic destined to the subscriber using reverse OIF mapping enabled on a specified multicast interface. Reverse OIF mapping is used to determine the subscriber VLAN interface and the multicast traffic bandwidth on the interface.

Levels of Shaping-Rate Adjustments

Both absolute and delta adjustments are made to a subscriber's aggregate shaping rate on a level 3 scheduler node.

 Copyright © 2013, Juniper Networks, Inc.

Adjustments that occur on the scheduler node can also impact the shaping rates for all queues. This adjustment can be undesirable for service providers who want to provide a premium level of service on specific queues.

For delta-based adjustments by multicast applications, you can control the distribution of shaping rates among queues by assigning the percentage of adjustment allowed for each queue. In addition, you can set a minimum adjusted shaping rate for each queue.

Figure 34 on page 1081 shows a sample multicast network with shaping rates adjusted at the scheduler node level. The shaping rate is reduced by 4 Mbps (from 41 Mbps to 37 Mbps) at the scheduler node for subscriber interface 1, which reduces the rates of both the best effort and video on demand (VoD) service queues.

Figure 34: Scheduler Node and Queues with Adjusted Shaping Rates

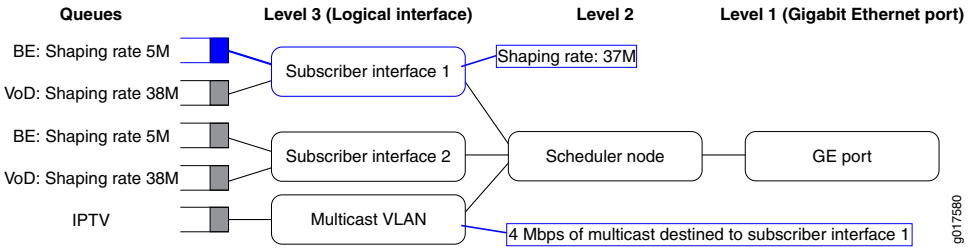

Figure 35 on page 1081 shows the same network with queue-based adjustments enabled for the best-effort queue on subscriber 1. The shaping rate of the best-effort queue is reduced by 4 Mbps (from 5 Mbps to 1 Mbps). The VoD service queue is not affected.

Figure 35: Queue with Adjusted Shaping Rate

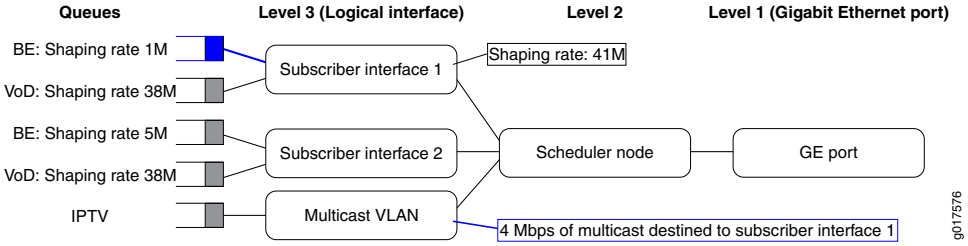

Related Documentation

- Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers on page 1096
- Configuring Shaping-Rate Adjustments on Queues on page 1097
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 1082
- Disabling Hierarchical Bandwidth Adjustment for Subscriber Interfaces with Reverse-OIF Mapping on page 1105
- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 1114

Copyright © 2013, Juniper Networks, Inc.

Shaping Rate Adjustments for Subscriber Local Loops Overview

This overview describes how an MX Series 3D Universal Edge Router installed as an edge router can adjust hierarchical CoS policy for subscriber interfaces for subscriber local loops. You can configure the router to throttle the traffic sent to subscriber local loops so that the traffic does not exceed the current data transmission rate of those lines. This feature ensures that changes to subscriber local loop speeds do not cause bandwidth contention at the subscriber's residential gateway.

In a typical subscriber access network, traffic destined to a subscriber is delivered from the access network, through an edge router, to a DSLAM. The DSLAM multiplexes subscriber traffic through a DSL, also known as a *local loop*, to the subscriber's residential gateway. When line noise or cross talk in a subcarrier causes the error rate on a DSL to exceed a certain threshold, the DSLAM can adapt itself by lowering the data transmission rate to that carrier device. A lower data transmission rate is less susceptible to induced errors.

You can configure an MX Series router to adjust the configured shaping rates on scheduler nodes for subscriber interfaces that represent subscriber local loops. Whenever a DSLAM resynchronizes a subscriber local loop speed, the router adjusts the configured shaping rate for that line so that the aggregate egress traffic to those subscribers is shaped to the local loop speed before the traffic reaches the DSLAM. Unless the maximum amount of bandwidth allocated to the subscriber interface on the router is throttled to the local loop speed, bandwidth contention can occur at the subscriber's residential gateway, which can cause the DSLAM to drop packets. This type of shaping-rate adjustment requires the topology discovery and traffic-monitoring features of the Access Node Control Protocol (ANCP).

You can configure ANCP to communicate the subscriber local loop speed to the MX Series router, which in turn throttles traffic destined to the associated subscriber interface so that it matches the subscriber local loop speed. ANCP acquires subscriber line rate information from DSLAMs and then communicates this data transmission rate for use with CoS.

Related Documentation

- Hierarchical CoS Shaping-Rate Adjustments Overview on page 1080
- Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 1083
- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1099
- Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1104
- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 1114
- For more information about the ANCP protocol, see the ANCP Topology Discovery and Traffic Reporting Overview on page 1311.

Copyright © 2013, Juniper Networks, Inc.

Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops

These guidelines apply to configuring an MX Series 3D Universal Edge Router installed as an edge router to adjust the configured shaping rates on scheduler nodes for subscriber interfaces that represent subscriber local loops. This shaping-rate feature uses the topology discovery and traffic-monitoring features of ANCP.

When you enhance hierarchical CoS policy by configuring ANCP-driven shaping-rate adjustments, consider the following guidelines:

- Shaping-rate adjustments are supported on EQ DPCs and MPCs on MX Series routers.
- Shaping-rate adjustments are supported only for subscriber local loops that terminate at DSLAMs that you have configured as ANCP neighbors of the MX Series router.
- Shaping-rate adjustments are supported only for scheduler nodes for which you have configured an initial shaping rate by including the **shaping-rate** statement in a traffic-control profile applied to the scheduler node. Specify the initial shaping rate as a peak rate, in bits per second (bps), and not as a percentage. Other methods of configuring a shaping rate are not supported with this feature.
- Shaping-rate adjustments are supported only for scheduler nodes that are static logical interface sets that you have configured to operate at Level 3 of the scheduler hierarchy on the router. If an interface set is configured with a logical interface (such as unit 0) and queue, then the interface set is an internal scheduler node (as opposed to a root node or a leaf node) at Level 2 of the hierarchy. However, if there are no traffic control profiles are configured on logical interfaces in an interface set, then the interface set is an internal scheduler node at Level 3 of the hierarchy.
- Shaping-rate adjustments are supported only for subscriber interfaces over physical interfaces that you have configured to operate in hierarchical scheduler mode. Only ports on EQ DPCs in MX Series routers support hierarchical scheduler mode.
- After shaping-rate adjustments are enabled and the router has performed shaping-rate adjustments on a scheduler node, you can configure a new shaping rate by including the **shaping-rate** statement in a traffic-control profile and then applying that profile to that scheduler node. However, this new shaping-rate value does not immediately result in shaping traffic at the new rate. The scheduler node continues to be shaped at rate set by ANCP. Only when the ANCP shaping-rate adjustment feature is disabled is the scheduler node shaped at the newly configured shaping-rate.
- The Layer 2 Tunneling Protocol (L2TP) is often used to carry traffic securely between an L2TP Network Server (LNS) and an L2TP Access Concentrator (LAC). The QoS adjustment feature supports the shaping overhead options that you can use to add a specified number of bytes to the actual packet length when determining shaped session packet length. ANCP shaping-rate adjustments are not supported for ingress traffic, only for egress traffic. To configure the number of bytes to add to the packet at the egress side of the tunnel, include the **egress-shaping-overhead** and **mode** statements at the **[edit chassis fpc** *slot-number* **pic** *pic-number* **traffic-manager]** hierarchy level. Use the shaping overhead options if you need to account for encapsulation overhead.

For more information about the ANCP protocol, see the "ANCP Topology Discovery and Traffic Reporting Overview" on page 1311.

Related Documentation

- Hierarchical CoS Shaping-Rate Adjustments Overview on page 1080
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 1082
- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1099
- Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1104
- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 1114

CoS Adjustment Control Profiles Overview

CoS adjustment control profiles control which applications and algorithms can modify a subscriber's shaping characteristics after a subscriber is instantiated. Subscriber shaping characteristics are configured using the Junos OS CLI or by RADIUS messages. Adjustment control profiles enable subscriber shaping characteristics by to be adjusted by other applications like ANCP, PPPoE tags, and RADIUS Change of Authorization (CoA) after a subscriber is instantiated. Adjustment control profiles are router-wide and apply to both static and dynamic interfaces.

Table 99 on page 1084 describes the terms used to define shaping characteristics.

Table 99: Terms and Definitions for Shaping Characteristics

Term	Definition
Shaping-rate	The maximum rate of a scheduler node or queue. Also known as Peak Information Rate (PIR).
Overhead-accounting mode	A class-of-service traffic-control profile attribute that specifies whether the downstream network is a frame-based network, like Ethernet, or a cell-based network, like ATM. NOTE: The downstream network is not necessarily the directly attached network. In typical broadband services router (BSR) network configurations, the directly attached network is an Ethernet access network, which provides access to either another frame-based network, or a cell-based network. When cell-mode is specified, the Juniper Networks router adjusts rates (like the shaping-rate) to "rate * 48/53" to account for 5-byte ATM AAL5 headers.
Overhead-accounting bytes	A class-of-service traffic-control profile attribute that specifies the number of bytes per packet to be included or excluded from the shaping mechanism. For example, to properly account for a 4-byte header stripped by the downstream network, set the overhead-accounting bytes to -4. To properly account for a 12-byte header added by the downstream network, set the overhead-accounting bytes to 12.
Effective-shaping-rate	The shaping-rate resulting from three attributes: shaping-rate, overhead-accounting mode, and overhead-accounting bytes.
Rate Adjustment	The adjustment that other applications like ANCP, PPPoE tagging, and RADIUS CoA can make when the shaping rate is configured through the CLI or RADIUS.

 Copyright © 2013, Juniper Networks, Inc.

Effective Shaping Rate

CoS is responsible for communicating shaping rate information from the Routing Engine to the Packet Forwarding Engine. The shaping-rate, also known as peak information rate (PIR), is the maximum rate for a scheduler node or queue.

The true rate of a subscriber at the access-loop/CPE is a function of:

- The shaping-rate in effect for the subscriber's household, in bits per second.
- Whether the subscriber is connected to a frame-based or cell-based network. This attribute is known in CoS as the overhead-accounting mode.
  - **overhead-accounting frame-mode-bytes**—When the **overhead-accounting** mode is set to **frame-mode**, this is the number of bytes in each frame that are accounted for by the shaper. This value represents the number of bytes that are encapsulated and decapsulated by the downstream equipment.
  - **overhead-accounting cell-mode-bytes**—When overhead-accounting mode is set to **cell-mode**, this is the number of bytes in each frame that are accounted for by the shaper for a downstream cell-mode network. This value represents the number of bytes that are encapsulated and decapsulated by the downstream equipment.

You configure the values for the **shaping-rate** and **overhead-accounting** mode options under either the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level or the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level. These options are supported on MX Series routers. The applications and specified algorithms configured in the adjustment control profile use the values of these options to adjust the shaping rate for static and dynamically instantiated subscribers.

NOTE: Chassis *egress-shaping-overhead* is not included in the effective rate. Egress-shaping-overhead accounts for the physical interface overhead (ISO OSI Layer 1). Effective shaping-rate is a Layer 2 (ISO OSI) rate.

Applications and Associated Algorithms in Adjustment Control Profiles

Table 100 on page 1085 describes the applications and their associated default algorithms that can be configured to perform rate adjustments after the subscriber is instantiated.

Table 100: Adjustment Control Profile Applications and Algorithms

Application	Default Priority	Default Algorithm	Description
RADIUS-CoA	1	Adjust-always	RADIUS CoA messages can update the subscriber's attributes (like shaping rate) after the subscriber is authenticated and QoS parameters (like shaping rate) are assigned.

Table 100: Adjustment Control Profile Applications and Algorithms *(continued)*

Application	Default Priority	Default Algorithm	Description
ANCP	1	Adjust-always	The ANCP application can modify the existing shaping rate for both static and dynamic logical interfaces, and static interface sets. By default, ANCP can override all other applications. The shaping rate must be specified in order to override it.
PPPoE-Tags	2	Adjust-less	The PPPoE IA tag access-rate-downstream can modify the Junos OS CLI configured shaping- rate value, as well as the RADIUS shaping- rate value. By default, these values can be modified by subsequent RADIUS CoA messages and ANCP actions. These values are conveyed in PPPoE (PADI) discovery packets.

NOTE: **The lower the priority value, the higher the priority.**

You must enable each application to perform rate adjustments. Rate adjustments are global and affect all static and dynamically instantiated subscribers. The following rules apply to adjustment control profiles:

- If no adjustment control profile is configured, the default adjustment control profile is used.
- You can configure a maximum of one adjustment control profile; a commit error occurs if you configure more than one adjustment control profile.
- If an application is not configured with an adjustment control profile, Junos OS uses its default values for priority and algorithm. For example, if ANCP is not configured in the adjustment control profile, the ANCP application is set to a priority of 1 and the algorithm is set to adjust-always.
- Adjustment control profiles apply to both static and dynamic interfaces.
- You can configure the algorithm to the following values:
  - Adjust-never
  - Adjust-always
  - Adjust less
  - Adjust less than or equal
  - Adjust greater
  - Adjust greater than or equal
- When you modify an adjustment control profile, the changes take effect immediately and the modified profile is used for all further adjustments. However, existing adjustments are not reevaluated when you modify the adjustment control profile.

Copyright © 2013, Juniper Networks, Inc.

For example, if you have an ANCP adjustment that overrides a PPPoE adjustment on interface ge-1/1/0.100, and then you use the adjustment control profile to change the priority so that the ANCP priority is now lower than the PPPoE priority, Junos OS does not go back and reevaluate the adjustment on ge-1/1/0.100.

Related Documentation

- Configuring CoS Adjustment Control Profiles on page 1106
- Verifying the CoS Adjustment Control Profile Configuration on page 1107
- **adjustment-control-profiles on page 1442**

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 60

Configuring Bandwidth Management Parameters for Dynamic CoS

- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 1090
- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 1092
- Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces on page 1093
- Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces on page 1093
- Verifying the Number of Dedicated Queues Configured on MIC and MPC Interfaces on page 1095
- Providing Accurate Scheduling for a Demux Subscriber Interface of Aggregated Ethernet Links on page 1095
- Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers on page 1096
- Configuring Shaping-Rate Adjustments on Queues on page 1097
- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1099
- Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1104
- Disabling Hierarchical Bandwidth Adjustment for Subscriber Interfaces with Reverse-OIF Mapping on page 1105
- Verifying the Configuration of Shaping-Rate Adjustments for Subscriber Local Loops on page 1105
- Verifying the Configuration of ANCP for Shaping-Rate Adjustments on page 1106
- Configuring CoS Adjustment Control Profiles on page 1106
- Verifying the CoS Adjustment Control Profile Configuration on page 1107
- Reporting the Effective Shaping Rate for Subscribers on page 1108
- Verifying the Effective Shaping Rate Reporting Configuration on page 1108

Copyright © 2013, Juniper Networks, Inc.

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces

Service providers often used tiered services that must utilize excess bandwidth as traffic patterns vary. By default, excess bandwidth between a configured guaranteed rate and shaping rate is shared equally among all queues with the same excess priority value, which might not be optimal for all subscribers to a service.

This feature is supported for MIC and MPC interfaces on MX Series routers.

To configure parameters to manage excess bandwidth for subscriber interfaces:

1. Configure the parameters for the interface.
   a. Configure the guaranteed and shaping rates.
      i. Configure the guaranteed rate:

         [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]
         user@host# **set** guaranteed-rate(*rate* | $junos-cos-guaranteed-rate) <burst-size (*bytes* | $junos-cos-guaranteed-rate-burst)>

      ii. Configure the shaping rate:

         [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]
         user@host# **set** shaping-rate (*rate* | $junos-cos-shaping-rate) <burst-size (*bytes* | $junos-cos-shaping-rate-burst)>

      TIP: On MPC/MIC interfaces, the guaranteed rate and the shaping rate share the value specified for the burst size. If the guaranteed rate has a burst size specified, it is used for the shaping rate; if the shaping rate has a burst size specified, it is used for the guaranteed rate. If you have specified a burst for both rates, the system uses the lesser of the two values.

   b. Configure a rate for excess bandwidth.

      You can configure an excess rate for all priorities of traffic:

      [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]
      user@host# **set** excess-rate (percent *percentage* | $junos-cos-excess-rate) | proportion *value*)

      Optionally, you can configure an excess rate specifically for high- and low-priority traffic. When you configure the **excess-rate** statement for an interface, you cannot also configure the **excess-rate-low** and **excess-rate-high** statements.

      [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]
      user@host# **set** excess-rate-high(percent *percentage* | $junos-cos-excess-rate-high) | proportion *value*)

 Copyright © 2013, Juniper Networks, Inc.

```
user@host# set excess-rate-low (percent percentage | $junos-cos-excess-rate-low)
  | proportion value )
```

BEST PRACTICE: We recommend that you configure either a percentage or a proportion of the excess bandwidth for all schedulers with the same parent in the hierarchy. For example, if you configure interface 1.1 with twenty percent of the excess bandwidth, configure interface 1.2 with eighty percent of the excess bandwidth.

2. (Optional) Configure parameters for the queue.

   a. Configure the shaping rate.

   ```
   [edit dynamic-profiles profile-name class-of-service scheduler scheduler-name]
   user@host#set shaping-rate (rate | $junos-cos-scheduler-shaping-rate) <burst-size
     bytes>
   ```

   b. Configure the excess rate.

   ```
   [edit dynamic-profiles profile-name class-of-service scheduler scheduler-name]
   user@host#set excess-rate (percent percentage | percent
     $junos-cos-scheduler-excess-rate)
   ```

   c. (Optional) Configure the priority of excess bandwidth for the queue.

   ```
   [edit dynamic-profiles profile-name class-of-service scheduler scheduler-name]
   user@host#set excess-priority (low | high | $junos-cos-scheduler-excess-priority
     | none)
   ```

   TIP:

   For queues, you cannot configure the excess rate or excess priority in these cases:

   - When the **transmit-rate exact** statement is configured. In this case, the shaping rate is equal to the transmit rate and the queue does not operate in the excess region.
   - When the scheduling priority is configured as **strict-high**. In this case, the queue gets all available bandwidth and never operates in the excess region.

   By default, when traffic exceeds the shaping or guaranteed rates, the system demotes traffic configured with high or medium priority. To disable priority demotion, specify the **none** option. You cannot configure this option for queues configured with **transmit-rate** expressed as a percent and when the parent's guaranteed rate is set to zero.

Related Documentation

- For hardware requirements and configuration guidelines, see Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950

Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates

You can configure the overhead accounting feature to shape downstream traffic based on either frames or cells.

You can also account for the different byte sizes per encapsulation by configuring a byte adjustment value for the shaping mode.

This feature is supported on MPCs on MX Series routers.

To configure the overhead accounting feature in a dynamic profile:

1. Do one of the following to configure the shaping mode:
   - Specify the shaping mode.

     Frame shaping mode is enabled by default.

     [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*
     user@host#set overhead-accounting **(frame-mode | cell-mode)**

   - Configure a variable for the shaping mode.

     [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*
     user@host#set overhead-accounting **$junos-cos-shaping-mode**

2. (Optional) Do one of the following to configure the byte adjustment value:
   - Specify a byte adjustment value.

     [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*
     user@host#set overhead-accounting **bytes *byte-value***

   - Configure a variable for the byte adjustment.

     [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*
     user@host#set overhead-accounting **bytes $junos-cos-byte-adjust**

     BEST PRACTICE: We recommend that you specify a byte adjustment value that represents the difference between the customer premise equipment (CPE) protocol overhead and B-RAS protocol overhead.

     The available range is –120 through 124 bytes. The system rounds up the byte adjustment value to the nearest multiple of 4. For example, a value of 6 is rounded to 8, and a value of -10 is rounded to -8.

Related Documentation

- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072
- Example: Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 1111

Copyright © 2013, Juniper Networks, Inc.

• Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 986

Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces

To configure the PPPoE vendor-specific tags feature in a dynamic profile:

NOTE: When you enable this feature, the values supplied by the PPPoE vendor-specific tags override the parameters that you have configured for shaping-rate and overhead-accounting statements at the **[edit dynamic-profiles profile-name class-of-service traffic-control-profile]** hierarchy level.

1. (Optional) To configure the shaping rate based on access line information:

   [edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]
   user@host# **set vendor-specific-tags actual-data-rate-downstream**

2. (Optional)To configure the overhead-accounting based on access-line information:

   [edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]
   user@host# **set vendor-specific-tags access-loop-encapsulation**

Related Documentation

• Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 1074
• Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072

Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces

This topic describes how to manage dedicated and remaining queues for static and dynamic subscriber interfaces configured in dynamic profiles.

You manage queues at the chassis and physical port level in the static configuration hierarchies, then configure dynamic scheduling and shaping parameters for the subscriber interfaces in the dynamic profile.

• Configuring the Maximum Number of Queues for MIC and MPC Interfaces on page 1093
• Configuring Remaining Common Queues on MIC and MPC Interfaces on page 1094

Configuring the Maximum Number of Queues for MIC and MPC Interfaces

30-Gigabit Ethernet Queuing MPCs and 60-Gigabit Ethernet Queuing and Enhanced Queuing MPCs support a dedicated number of queues when configured for hierarchical scheduling and per-unit scheduling configurations.

To scale the number of subscriber interfaces per queue, you can modify the number of queues supported on the MIC.

To configure the number of queues:

1. Specify that you want to configure the MIC.

   user@host# **edit chassis fpc** *slot-number* **pic** *pic-number*

2. Configure the number of queues.

   [edit chassis fpc *slot-number* pic *pic-number*]
   user@host# **setmax-queues-per-interface (8 | 4)**

Configuring Remaining Common Queues on MIC and MPC Interfaces

30-Gigabit Ethernet Queuing MPCs and 60-Gigabit Ethernet Queuing and Enhanced Queuing MPCs support a dedicated set of queues when configured with hierarchical scheduling.

When the number of dedicated queues is reached on the module, there can be queues remaining. Traffic from these logical interfaces are considered unclassified and attached to a common set of queues that are shared by all subsequent logical interfaces.

You can configure traffic shaping and scheduling resources for the remaining queues by attaching a special traffic-control profile to the interface. This feature enables you to provide the same shaping and scheduling to remaining queues as the dedicated queues.

To configure the remaining queues on a MIC or MPC interface:

1. Configure CoS parameters in a traffic-control profile.

   [edit class-of-service]
   user@host# **edit traffic-control-profiles** *profile-name*

2. Enable hierarchical scheduling for the interface.

   [edit interfaces *interface-name*]
   user@host# **set** hierarchical-scheduler

3. Attach the traffic control profiles for the dedicated and remaining queues to the port on which you enabled hierarchical scheduling.

   To provide the same shaping and scheduling parameters to dedicated and remaining queues, reference the same traffic-control profile.

   a. Attach the traffic-control profile for the dedicated queues on the interface.

      [edit class-of-service interfaces *interface-name*]
      user@host# **set output-traffic-control-profile** *profile-name*

   b. Attach the traffic-control profile for the remaining queues on the interface.

      [edit class-of-service interfaces *interface-name*]
      user@host# **set output-traffic-control-profile-remaining** *profile-name*

Related Documentation

- Verifying the Number of Dedicated Queues Configured on MIC and MPC Interfaces on page 1095
- Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview on page 1076

 Copyright © 2013, Juniper Networks, Inc.

- Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 963
- Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 965

Verifying the Number of Dedicated Queues Configured on MIC and MPC Interfaces

Purpose Display the number of dedicated queue resources that are configured for the logical interfaces on a port.

Action

```
user@host#show class-of-service interface ge-1/1/0
Physical interface: ge-1/1/0, Index: 166
Queues supported: 4, Queues in use: 4
Total non-default queues created: 4
  Scheduler map: <default>, Index: 2
  Chassis scheduler map: <default-chassis>, Index: 4

  Logical interface: ge-1/1/0.100, Index: 72, Dedicated Queues: no
    Shaping rate: 32000
    Object              Name                   Type                   Index
    Scheduler-map       <remaining>                                   0
    Classifier          ipprec-compatibility   ip                     13

  Logical interface: ge-1/1/0.101, Index: 73, Dedicated Queues: no
    Shaping rate: 32000
    Object              Name                   Type                   Index
    Scheduler-map       <remaining>                                   0
    Classifier          ipprec-compatibility   ip                     13

  Logical interface: ge-1/1/0.102, Index: 74, Dedicated Queues: yes
    Shaping rate: 32000
    Object                   Name               Type                  Index
    Traffic-control-profile  <control_tc_prof>  Output                45866
```

Related Documentation

- *Managing Dedicated and Remaining Queues for Static CoS Configurations on MIC and MPC Interfaces*
- Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces on page 1093

Providing Accurate Scheduling for a Demux Subscriber Interface of Aggregated Ethernet Links

Unlike VLAN subscriber interfaces, enabling link protection is not required for configuring hierarchical CoS on demux interfaces. Instead, we recommend that you enable targeted distribution on the demux interface to provide accurate scheduling for the aggregated Ethernet links.

Before you begin, configure the subscriber interface with aggregated Ethernet:

- For static and dynamic IP demux interfaces, see "Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet" on page 825.

- For static and dynamic VLAN demux interfaces, see "Configuring a Static or Dynamic VLAN Demux Subscriber Interface over Aggregated Ethernet" on page 826.

To provide accurate scheduling for a demux subscriber interface of aggregated Ethernet links:

1. Enable targeted distribution for the demux interface.

   See "Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces" on page 828.

2. Enable hierarchical scheduling on the link aggregation bundle.

   See "Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links" on page 1023.

3. (Optional) Enable module redundancy to ensure that CoS resources are provisioned for the aggregated Ethernet links if a module or a link fails. By default, link redundancy is supported.

   See "Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface" on page 829.

4. (Optional) Configure rebalancing periodically or manually for the subscribers. See "Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface" on page 829.

5. Attach static or dynamic traffic shaping and scheduling parameters at the aggregated Ethernet logical interface or its underlying physical interface. See:

   - Configuring Traffic Scheduling and Shaping for Subscriber Access on page 971
   - Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973
   - Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile on page 979
   - Applying Minimal Shaping and Scheduling to Remaining Subscriber Traffic on page 980

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 830

Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers

- Overview on page 1097
- Configuring a Static Minimum Adjusted Shaping Rate on Scheduler Nodes on page 1097
- Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes on page 1097

 Copyright © 2013, Juniper Networks, Inc.

Overview

Absolute adjustments and delta adjustments are performed at the scheduler node level. You can configure a minimum adjusted shaping rate at the scheduler node level using static or dynamic CoS parameters.

This feature is supported for adjustments performed by the ANCP and multicast applications on both EQ DPCs and MPC/MIC modules on MX Series routers.

BEST PRACTICE: **For multicast traffic, you can configure a minimum adjusted shaping rate at the queue level. We recommend that you configure the minimum adjusted value at the scheduler node or the queue, but not both.**

When you configure a minimum adjusted value for a node and for a scheduler that is referenced by a scheduler map in the same traffic-control-profile, the system uses the minimum value from the scheduler.

This feature is supported for adjustments performed by the ANCP and multicast applications on both EQ DPCs and MPC/MIC modules on MX Series routers.

Configuring a Static Minimum Adjusted Shaping Rate on Scheduler Nodes

To apply a minimum adjusted shaping rate for a scheduler node:

- Configure the **adjust-minimum** statement for the static traffic-control profile.

  [edit class-of-service traffic-control-profiles *profile-name*]
  user@host# **set adjust-minimum** *rate*

Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes

To apply a minimum adjusted shaping rate for a scheduler node:

- Configure the **adjust-minimum** statement for the dynamic traffic-control profile.

  [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]
  user@host# **set** adjust-minimum *rate*

Related Documentation

- Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 986
- Configuring Shaping-Rate Adjustments on Queues on page 1097
- Hierarchical CoS Shaping-Rate Adjustments Overview on page 1080

Configuring Shaping-Rate Adjustments on Queues

- Overview on page 1098
- Configuring a Static Shaping-Rate Adjustment for Queues on page 1098
- Configuring a Dynamic Shaping-Rate Adjustment for Queues on page 1099

Overview

By default, the multicast application adjusts the shaping rates at the scheduler node level. This adjustment also impacts the shaping rates for all queues, which can be undesirable for service providers who want to provide a premium level of service on specific queues.

For multicast applications, you can control the distribution of shaping rates among queues by assigning the percentage of adjustment allowed for each queue. In addition, you can set a minimum adjusted shaping rate for each queue.

This feature is supported for adjustments performed by the multicast application on MPC/MIC modules on MX Series routers.

BEST PRACTICE: We recommend that you configure the minimum adjusted value at the scheduler node or the queue, but not both.

When you configure a minimum adjusted value for a node and for a scheduler that is referenced by a scheduler map in the same traffic-control-profile, the system uses the minimum value from the scheduler.

This feature is supported for adjustments performed by the multicast application on MPC/MIC modules on MX Series routers.

Configuring a Static Shaping-Rate Adjustment for Queues

To configure adjustment parameters for a queue:

1. Configure the percentage of adjustment for the shaping rate.

   [edit class-of-service schedulers *scheduler-name*]
   user@host# **set adjust-percent** *percentage*

2. Configure the minimum adjusted value for the shaping rate.

   Do one of the following:

   - Configure the minimum adjusted value for the queue.

     [edit class-of-service schedulers *scheduler-name*]
     user@host# **set adjust-minimum** *rate*

   - Configure the minimum adjusted value for the node.

     [edit class-of-service traffic-control-profile *profile-name*]
     user@host# **set adjust-minimum** *rate*

BEST PRACTICE: Ensure that the minimum adjusted value that you configure does not exceed the shaping rate and is not lower than the configured transmit rate.

Copyright © 2013, Juniper Networks, Inc.

Configuring a Dynamic Shaping-Rate Adjustment for Queues

To configure adjustment parameters for a queue in a dynamic profile:

1. Configure the percentage of adjustment for the shaping rate.

   [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
   user@host# **set** adjust-percent ***percentage***

2. Configure the minimum adjusted value for the shaping rate.

   Do one of the following:

   - Configure the minimum adjusted value for the queue.

     [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
     user@host# **set** adjust-minimum (*rate* | $junos-cos-adjust-minimum)

   - Configure the minimum adjusted value for the node.

     [edit dynamic-profiles *profile-name* class-of-service traffic-control-profile *profile-name*]
     user@host# **set** adjust-minimum *rate*

BEST PRACTICE: **Ensure that the minimum adjusted value that you configure does not exceed the shaping rate and is not lower than the configured transmit rate.**

Related Documentation

- Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 986
- Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers on page 1096
- Hierarchical CoS Shaping-Rate Adjustments Overview on page 1080

Enabling Shaping-Rate Adjustments for Subscriber Local Loops

You can enhance a CoS implementation by enabling an MX Series 3D Universal Edge Router to adjust the hierarchical CoS policy shaping rate configured for static interface sets that consist of two or more VLANs and represent subscriber local loops. Whenever the digital subscriber line access multiplexer (DSLAM) resynchronizes its data transmission rate to a digital subscriber line (DSL), the router adjusts the shaping rate for the associated subscriber interface so that the maximum bandwidth allocation cannot exceed the current data rate for the associated subscriber local loop. This feature ensures that data

transmission rate adjustments by the DSLAM do not cause bandwidth contention at the subscriber's residential gateway.

This topic includes the following tasks:

- Configuring Static Logical Interface Sets to Serve as CoS Hierarchical Scheduler Nodes for Subscriber Loops on page 1100
- Configuring the Logical Interfaces That Compose the Static Logical Interface Sets on page 1101
- Configuring Hierarchical CoS on the Static Logical Interface Sets That Serve as Hierarchical Scheduler Nodes for Subscriber Local Loops on page 1101
- Configuring ANCP Functionality That Supports and Drives Shaping-Rate Adjustments for Subscriber Local Loops on page 1103

Configuring Static Logical Interface Sets to Serve as CoS Hierarchical Scheduler Nodes for Subscriber Loops

To configure a logical interface set, begin by including the **interface-set** statement with the *interface-set-name* option at the **[edit interfaces]** hierarchy level.

An interface set is composed of two or more logical interfaces on the same physical interface. Each logical interface in an interface set corresponds to an individual subscriber service, such as voice, video, or data. To specify either a list of logical unit numbers or the single outer VLAN tag used to identify the logical interfaces that compose the interface set, include statements at the **[edit interfaces interface-set *interface-set-name*]** hierarchy level:

- For an interface set composed of a list of logical interfaces identified by an inner VLAN tag on Ethernet frames (called the customer VLAN, or C-VLAN, tag), you must specify each logical interface by including the **unit** statement with the ***logical-unit-number*** option.

```
[edit]
interfaces {
   interface-set interface-set-name {
      interface ethernet-interface-name { # EQ DPC port
         unit logical-unit-number;
         unit logical-unit-number;
         . . .
      }
      . . .
   }
}
```

- For an interface set composed of a set of VLANs grouped at the DSLAM and identified by the same service VLAN (S-VLAN) tag), you must specify the S-VLAN tag as the outer VLAN tag for each VLAN by including the **vlan-tags-outer** statement with the ***vlan-tag*** option.

```
[edit]
interfaces {
   interface-set interface-set-name {
      interface ethernet-interface-name { # EQ DPC port
         vlan-tags-outer vlan-tag; # Identify the DSLAM
```

 Copyright © 2013, Juniper Networks, Inc.

```
        }
        ...
      }
    }
```

For more information about configuring CoS hierarchical schedulers, see the *Junos OS Class of Service Library for Routing Devices*.

Configuring the Logical Interfaces That Compose the Static Logical Interface Sets

Each underlying physical interface must be configured to operate in hierarchical scheduler mode and to support stacked VLAN tagging on all logical interfaces. To configure, include the **hierarchical-scheduler** statement and the **stacked-vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level.

To associate the individual logical interfaces of an interface set with specific subscriber services provided by the subscriber local loop, bind an S-VLAN tag and a C-VLAN tag to each logical interface that belongs to a scheduler node that represents a subscriber local loop. Ethernet frames sent from the logical interfaces contain an outer VLAN tag that identifies a DSLAM and an inner VLAN tag that identifies a subscriber port on the DSLAM. To configure, include the **vlan-tags** statement at each logical interface:

```
[edit]
interfaces {
  ethernet-interface-name { # EQ DPC port underlying an interface set
    hierarchical-scheduler;
    stacked-vlan-tagging; # Support 802.1Q VLAN dual-tagged frames
    unit logical-unit-number { # Bind S-VLAN and C-VLAN tags to logical interface
      vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
    }
    ...
  }
}
```

For more information about configuring 802.1Q VLANs, see the *Junos OS Network Interfaces Library for Routing Devices*.

Configuring Hierarchical CoS on the Static Logical Interface Sets That Serve as Hierarchical Scheduler Nodes for Subscriber Local Loops

To configure hierarchical CoS on the static logical interface set that serves as the hierarchical scheduler node for a subscriber local loop:

1. For each scheduler node that represents a subscriber local loop, configure an initial shaping rate.

NOTE: The CoS shaping-rate feature is supported only for scheduler nodes with a configured shaping rate. The initial shaping rate must be configured by applying a traffic-control profile that includes the shaping-rate statement. Specify the initial shaping rate as a peak rate, in bits per second (bps), and not as a percentage. Other methods of configuring a shaping rate are not supported with this feature.

- To enable traffic heading downstream (from the router to the DSLAM) to be gathered into an interface set, include the **interface-set** statement and define the logical interface set name as the ***interface-set-name*** option at the **[edit class-of-service interfaces]** hierarchy level.
- To apply output traffic scheduling and shaping parameters at the logical interface set level (rather than at the logical unit level), include the **output-traffic-control-profile** statement and specify the name of a traffic-control profile as the ***profile-name*** option at the **[edit class-of-service interfaces interface-set *interface-set-name*]** hierarchy level.

To configure, include the following statements:

```
interfaces { # Configure interface-specific CoS for incoming packets
  interface-set interface-set-name { # Configure a hierarchical scheduler
    output-traffic-control-profile tc-profile-name; # Level 3 scheduler node
  }
  . . .
}
traffic-control-profiles { # Define traffic-control profiles
  tc-profile-name { # Specify a scheduler map and traffic-shaping parameters
    scheduler-map map-name;
    shaping-rate rate; # This is the "configured shaping rate"
    guaranteed-rate (percent percentage | rate);
    delay-buffer-rate (percent percentage | rate);
  }
  . . .
}
```

You can include the statements at the following hierarchy levels:

- **[edit class-of-service]**
- **[edit dynamic-profiles *profile-name* class-of-service]**

2. Configure the scheduler maps referenced in the traffic-control profiles applied to the interface sets, the schedulers referenced in those scheduler maps, and the drop profiles referenced in those schedulers.

   - A scheduler map establishes the traffic output queues (forwarding classes) for a scheduler node and associates each queue with a specific scheduler map.
   - A scheduler defines queue properties (transmit rate, buffer size, priority, and drop profile) that specify how traffic is treated in the output queue.
   - A drop profile specifies how aggressively the MX Series router drops packets that are managed by a particular scheduler by defining either a segmented or interpolated graph that maps output queue fullness to packet drop probability.

   To configure, include the statements at the static **[edit class-of-service]** hierarchy level:

```
[edit]
class-of-service {
  scheduler-maps { # Assign queuing characteristics to output queues
    map-name { # Map output queues to
```

 Copyright © 2013, Juniper Networks, Inc.

```
            forwarding-class class-name scheduler scheduler-name;
            forwarding-class class-name scheduler scheduler-name;
            ...
        }
        ...
    }
    schedulers { # Define queuing characteristics
        scheduler-name { # Specify queuing and buffer management
            transmit-rate transmit-rate-option;
            buffer-size buffer-size-option;
            priority priority-level;
            drop-profile-map loss-priority loss-priority-option protocol any drop-profile
                drop-profile-name;
            . . .
        }
    }
    drop-profiles { # Define random early detection (RED) for the delay buffer
        drop-profile–name { # Specify how to drop packets from an output queue
            drop-profile-name { # Map a queue fullness to a drop probability
                fill-level percentage drop-probability percentage; # Option 1: segmented
                fill-level percentage drop-probability percentage;
                . . .
            }
                interpolate { # Option 2: interpolated
                    drop-probability [ values ];
                    fill-level [ values ];
                }
            }
        . . .
    }
}
```

For more information about configuring scheduler maps, schedulers, and drop profiles, see the *Junos OS Class of Service Library for Routing Devices*.

Configuring ANCP Functionality That Supports and Drives Shaping-Rate Adjustments for Subscriber Local Loops

To configure the Access Node Control Protocol (ANCP) functionality that supports and drives the shaping-rate adjustments for subscriber local loops:

- Enable ANCP to monitor subscriber local loop rates at the DSLAMs and communicate this information to CoS.
- Configure each DSLAM as an ANCP neighbor of the router so that TCP connections can be established between the router and each DSLAM.
- Identify the subscriber interface sets whose traffic is monitored and shaped by ANCP, and associate those interface sets with the corresponding identifiers configured on the access node (DSLAM) to uniquely identify the subscriber local loops within the access network.

  ANCP uses this information to build a mapping of subscribers to subscriber interfaces. When ANCP receives port management messages from a DSLAM or other access

node, it uses the access identifier contained in the message to determine which hierarchical scheduler node corresponds to the subscriber.

To configure, include statements at the **[edit protocols ancp]** hierarchy level:

```
[edit]
protocols {
  ancp {
    qos-adjust; # Enable ANCP to monitor and adjust CoS shaping rates
    neighbor ip-address; # Configure each DSLAM as an ANCP neighbor
    ...
    interfaces { # Identify subscribers for which ANCP can adjust shaping rates
      interface-set {
        interface-set-name {
          access-identifier identifier-string; # DSLAM ID for the local loop
        }
      }
      ...
    }
    ...
  }
  ...
}
```

Related Documentation

- For hardware requirements and configuration guidelines, see Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 1083
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 1082
- Verifying the Configuration of ANCP for Shaping-Rate Adjustments on page 1106
- Verifying the Configuration of Shaping-Rate Adjustments for Subscriber Local Loops on page 1105
- Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1104
- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 1114

Disabling Shaping-Rate Adjustments for Subscriber Local Loops

To disable hierarchical CoS shaping-rate adjustments for subscriber local loops:

- Disable hierarchical CoS traffic-shaping adjustment by ANCP:

```
[edit protocols ancp]
user@host# delete qos-adjust
```

Traffic-shaping parameters for all subscriber local loops revert to their current configured values.

Related Documentation

- For hardware requirements and configuration guidelines, see Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 1083
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 1082

Copyright © 2013, Juniper Networks, Inc.

- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1099
- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 1114

Disabling Hierarchical Bandwidth Adjustment for Subscriber Interfaces with Reverse-OIF Mapping

You can disable hierarchical bandwidth adjustment for all subscriber interfaces with reverse OIF mapping enabled on a specified multicast interface. Reverse OIF mapping is used to determine the subscriber VLAN interface and the multicast traffic bandwidth on the interface.

To disable hierarchical bandwidth adjustment:

1. Specify that you want to access the subscriber interfaces with reverse-OIF mapping enabled.

   [edit routing-instances *routing-instance* routing-options multicast interface *interface-name*]
   user@host# **edit reverse-oif-mapping**

2. Disable hierarchical bandwidth adjustment for all subscriber interfaces on the interface.

   user@host# **set no-qos-adjust**

Related Documentation

- Hierarchical CoS Shaping-Rate Adjustments Overview on page 1080
- *Example: Configuring Multicast with Subscriber VLANs*

Verifying the Configuration of Shaping-Rate Adjustments for Subscriber Local Loops

Purpose Display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical CoS.

NOTE: After shaping-rate adjustments are enabled and the router has performed shaping-rate adjustments on a scheduler node, you can configure a new shaping rate by including the **shaping-rate** statement in a traffic-control profile and then applying that profile to that scheduler node. However, this new shaping-rate value does not immediately result in shaping traffic at the new rate. The scheduler node continues to be shaped at rate set by ANCP. Only when the ANCP shaping-rate adjustment feature is disabled is the scheduler node shaped at the newly configured shaping-rate.

Action Issue the ***show class-of-service interface-set*** operational command.

Related Documentation

- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1099

Verifying the Configuration of ANCP for Shaping-Rate Adjustments

Purpose Use to display or clear information about the ANCP configuration for shaping-rate adjustments.

Action

- To display ANCP neighbor information, issue the ***show ancp neighbor*** operational command.
- To clear ANCP neighbors, issue the ***clear ancp neighbor*** operational command.
- To display ANCP subscriber information, issue the ***show ancp subscriber*** operational command.
- To display ANCP class-of-service information, issue the ***show ancp cos*** operational command.

If ANCP is not yet enabled, the process starts when you commit a configuration that contains the **protocols ancp** stanza.

Related Documentation

- ANCP Topology Discovery and Traffic Reporting Overview on page 1311
- Configuring ANCP on page 1332

Configuring CoS Adjustment Control Profiles

To configure adjustment control profiles:

NOTE: **You can only configure one adjustment control profile.**

1. Configure the adjustment control profile name.

   [edit]
   user@host# **edit class-of-service adjustment-control-profiles** ***profile-name***

2. (Optional) Configure the adjustment controls for the Access Node Control Protocol (ANCP) application:

   [edit class-of-service adjustment-control-profiles *profile-name*]
   user@host# **set application ancp priority** ***priority*** **algorithm** ***algorithm***

3. (Optional) Configure the adjustment controls for the RADIUS CoA application:

   [edit class-of-service adjustment-control-profiles *profile-name*]
   user@host# **set application radius-coa priority** ***priority*** **algorithm** ***algorithm***

4. (Optional) Configure the adjustment controls for the PPPoE tags:

   [edit class-of-service adjustment-control-profiles *profile-name*]
   user@host# **set application pppoe-tags priority** ***priority*** **algorithm** ***algorithm***

5. (Optional) Verify your configuration.

   `user@host>` **show class-of-service adjustment-control-profiles**

 Copyright © 2013, Juniper Networks, Inc.

```
name: ANCP, priority: 1, algorithm: less;
name: RADIUS CoA, priority: 1, algorithm: always;
name: PPPoE IA tags, priority: 2, algorithm: less;
```

Related Documentation

- CoS Adjustment Control Profiles Overview on page 1084
- Verifying the CoS Adjustment Control Profile Configuration on page 1107
- **adjustment-control-profiles on page 1442**
- **overhead-accounting (Dynamic Traffic Shaping) on page 1841**

Verifying the CoS Adjustment Control Profile Configuration

Purpose View the class-of-service (CoS) adjustment control profile.

Action

- To display the CoS adjustment control profile:

  user@host> **show class-of-service adjustment-control-profile** ***profile-name***

```
user@host> show class-of-service adjustment-control-profile acp1
  name: ANCP, priority: 1, algorithm: less
  name: RADIUS CoA, priority: 1, algorithm: always
  name: PPPoE IA tags, priority: 2, algorithm: less

user@host>
```

Related Documentation

- CoS Adjustment Control Profiles Overview on page 1084
- Configuring CoS Adjustment Control Profiles on page 1106
- **adjustment-control-profiles on page 1442**
- **application (Adjustment Control Profiles) on page 1459**

Reporting the Effective Shaping Rate for Subscribers

The Effective-Shaping-Rate VSA [26–177] provides the best estimate for a subscriber's downstream traffic rate for accounting purposes. The VSA is included in RADIUS Acct-Start, Acct-Stop, and Interim-Acct messages. The reported rate is the rate enforced on the L3, L2, or L1 node according to local policy. The value of the VSA varies depending on your configuration:

- Actual rate—When effective shaping rate reporting is enabled.
- Advisory rate—When the advisory rate is configured and effective shaping rate reporting is not enabled.
- Port speed—When the advisory rate is not configured and effective shaping rate reporting is not enabled.

When you disable reporting, the VSA reports either the advisory rate or port speed for both existing subscribers and new subscribers that log in after reporting is disabled.

To enable reporting of the actual downstream traffic rate:

- Enable reporting.

```
[edit chassis]
user@host1# set effective-shaping-rate
```

NOTE: When the traffic control profile for the subscriber specifies cell-mode, the effective shaping rate does not account for cell padding according to the encapsulation type. The rate includes the 48/53 cell tax.

Related Documentation

- Verifying the Effective Shaping Rate Reporting Configuration on page 1108
- Hierarchical CoS Shaping-Rate Adjustments Overview on page 1080
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS on page 115

Verifying the Effective Shaping Rate Reporting Configuration

Purpose Verify whether reporting is enabled for the effective shaping rate. Display the effective shaping rate when reporting is enabled.

Action

- To display configuration information for effective shaping rate reporting:

```
[edit]
user@host# show chassis
effective-shaping-rate;
aggregated-devices {
```

 Copyright © 2013, Juniper Networks, Inc.

```
    ethernet {
      device-
    }
  }
```

- To display the effective shaping rate in kilobits per second when reporting is enabled:

```
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: 31000000
...
```

Related Documentation

- Reporting the Effective Shaping Rate for Subscribers on page 1108

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 61

Bandwidth Management for Dynamic CoS Examples

- Example: Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 1111
- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 1114

Example: Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates

This topic describes two scenarios for which you can configure dynamic shaping parameters to account for packet overhead in a downstream network.

The RADIUS administrator supplies the initial values on the RADIUS server, and the service activation is performed at subscriber login.

Figure 36 on page 1111 shows the sample network that the examples reference.

Figure 36: Sample Network Topology for Downstream Traffic

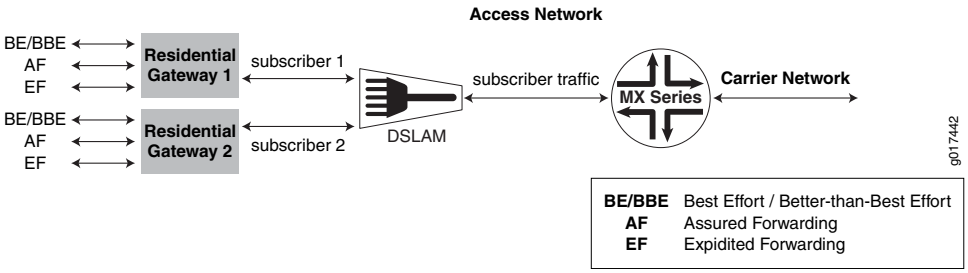

Managing Traffic with Different Encapsulations

In this example, the MX Series router shown in Figure 36 on page 1111 sends stacked VLAN frames to the DSLAM, and the DSLAM sends single-tagged VLAN frames to the residential gateway.

To accurately shape traffic at the residential gateway, the MX Series router must account for the different frame sizes. The difference between the stacked VLAN (S-VLAN) frames sent by the router and the single-tagged VLAN frames received at the residential gateway is a 4-byte VLAN tag. The residential gateway receives frames that are 4 bytes less.

Copyright © 2013, Juniper Networks, Inc.

To account for the different frame sizes, you configure the frame shaping mode with -4 byte adjustment:

1. Configure the traffic shaping parameters in the dynamic profile and attach them to the interface.

   Enabling the overhead accounting feature affects the resulting shaping rate, guaranteed rate, and excess rate parameters, if they are configured.

```
[edit]
dynamic-profiles {
  ethernet-downstream-network {
  interfaces {
    $junos-interface-ifd-name {
      unit $junos-underlying-interface-unit {
        family inet;
      }
    }
  }
  class-of-service {
    traffic-control-profiles {
      tcp-example-overhead-accounting-frame-mode {
        excess-rate percent $junos-cos-excess-rate
        guaranteed-rate $junos-cos-guaranteed-rate
        overhead-accounting $junos-cos-shaping-mode bytes $junos-cos-byte-adjust
        shaping-rate $junos-cos-shaping-rate;
      }
    }
    interfaces {
      $junos-interface-ifd-name {
        unit "$junos-underlying-interface-unit" {
        output-traffic-control-profile tcp1;
        }
      }
    }
  }
  }
}
```

   Table 101 on page 1112 lists the initial values defined by the RADIUS administrator for the shaping rates.

Table 101: Initial Shaping Values at Subscriber Login

Predefined Variable	RADIUS Tag	Value
$junos-cos-shaping-rate	T02	10m
$junos-cos-guaranteed-rate	T03	2m
$junos-cos-excess-rate	T05	50
$junos-cos-shaping-mode	T07	frame-mode
$junos-cos-byte-adjust	T08	–4

Copyright © 2013, Juniper Networks, Inc.

2. Verify the adjusted rates.

```
user@host#show class-of-service traffic-control-profile
Traffic control profile: tcp-example-overhead-accounting-frame-mode, Index:
61785
Excess rate 50
Shaping rate: 10000000
Guaranteed rate: 2000000
Overhead accounting mode: Frame Mode
Overhead bytes: -4
```

Managing Downstream Cell-Based Traffic

In this example, the DSLAM and residential gateway shown in Figure 36 on page 1111 are connected through an ATM cell-based network. The MX Series router sends Ethernet frames to the DSLAM, and the DSLAM sends ATM cells to the residential gateway.

To accurately shape traffic at the residential gateway, the MX Series router must account for the different physical network characteristics.

The administrator does not need to configure a byte adjustment value to account for the downstream ATM network, but has the option of configuring a byte adjustment value to account for different encapsulations or decapsulations.

To account for the different frame sizes, configure cell shaping mode:

1. Configure the traffic shaping parameters in the dynamic profile and attach them to the interface.

   Enabling the overhead accounting feature affects the resulting shaping rate, guaranteed rate, and excess rate parameters, if they are configured.

```
[edit]
dynamic-profiles {
  atm-downstream-network {
  interfaces {
    $junos-interface-ifd-name {
      unit $junos-underlying-interface-unit {
        family inet;
      }
    }
  }
  class-of-service {
    traffic-control-profiles {
      tcp-example-overhead-accounting-cell-mode {
        excess-rate percent $junos-cos-excess-rate
        guaranteed-rate $junos-cos-guaranteed-rate
        overhead-accounting $junos-cos-shaping-mode
        shaping-rate $junos-cos-shaping-rate
      }
    }
    interfaces {
      $junos-interface-ifd-name {
        unit "$junos-underlying-interface-unit" {
        output-traffic-control-profile tcp1;
        }
```

```
            }
          }
        }
        }
      }
```

Table 102 on page 1114 lists the initial values defined by the RADIUS administrator for the shaping rates.

Table 102: Initial Shaping Values at Subscriber Login

Predefined Variable	RADIUS Tag	Value
$junos-cos-shaping-rate	T02	10m
$junos-cos-guaranteed-rate	T03	2m
$junos-cos-excess-rate	T05	50
$junos-cos-shaping-mode	T07	cell-mode

2. Verify the adjusted rates.

```
user@host#show class-of-service traffic-control-profile
Traffic control profile: tcp-example-overhead-accounting-cell-mode, Index:
61785
Shaping rate: 10000000
Excess rate 50
Guaranteed rate: 2000000
Overhead accounting Cell Mode
Overhead bytes: 0
```

To account for ATM segmentation, the MX Series router adjusts all of the rates by 48/53 to account for ATM AAL5 encapsulation. In addition, the router accounts for cell padding, and internally adjusts each frame by 8 bytes to account for the ATM trailer.

Related Documentation

- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 1092

Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops

This example shows how you can enable shaping-rate adjustments for static logical interface sets that represent subscriber local loops:

1. Configure static logical interface sets to serve as CoS hierarchical scheduler nodes for subscriber local loops.

Copyright © 2013, Juniper Networks, Inc.

This example uses a single scheduler node that represents two subscriber local loops. The scheduler node is a static logical interface composed of two logical interfaces. The underlying physical interface is port **0** on a Gigabit Ethernet EQ DPC in slot 4, PIC 0:

```
[edit]
interfaces {
  interface-set ifset-of-logical-interfaces {
    interface ge-4/0/0 {
      unit 1;
      unit 2;
    }
  }
  ge-4/0/0 {
    description “access interface ge-4/0/0”;
    hierarchical-scheduler;
    stacked-vlan-tagging;
    unit 1 {
      description “DSL type ADSL1 = 0x01”;
      proxy-arp;
      vlan-tags outer 1 inner 1; # S-VLAN tag is ’1’ and C-VLAN tag is ’1’
      family inet { # Specify a secondary loopback address
        unnumbered-address lo0.0 preferred-source-address 192.168.7.3;
      }
    }
    unit 2 {
      description “DSL type ADSL1 = 0x01”;
      proxy-arp;
      vlan-tags outer 1 inner 2; # S-VLAN tag is ’1’ and C-VLAN tag is ’2’
      family inet { # Specify a secondary loopback address
        unnumbered-address lo0.0 preferred-source-address 192.168.7.4;
      }
    }
  }
}
```

2. Begin configuring hierarchical CoS on the static logical interface set that serves as the hierarchical scheduler node for the group of subscriber local loops.

```
[edit]
class-of-service {
  interfaces {
    interface-set ifset-of-logical-interfaces {
      output-traffic-control-profile tcp-premium-with-4–queues;
    }
  }
}
```

3. Configure the traffic-control profiles that can be applied to the scheduler node:

```
[edit]
class-of-service {
  traffic-control-profiles {
    tcp-basic-rate { # Specify a scheduler map and traffic controls
      shaping-rate 10m;
    }
    tcp-premium-with-4-queues { # Specify a scheduler map and traffic controls
```

```
            scheduler-map smap-premium-4q;
            shaping-rate 20m;
            guaranteed-rate 10m;
            delay-buffer-rate 5m;
        }
    }
}
```

In this example, the **tcp-premium-with-4-queues** traffic-control profile is applied to the interface set. The other profile provides a lower shaping rate and no guaranteed rate.

4. Configure the scheduler map **smap-premium-4q** that is referenced in the traffic-control profile for the scheduler node:

```
[edit]
class-of-service {
    scheduler-maps { # Define the queues that comprise each scheduler node
        smap-premium-4q { # Map each queue in the scheduler node to a scheduler
            forwarding-class be scheduler be_sch;
            forwarding-class af scheduler af_sch;
            forwarding-class ef scheduler ef_sch;
            forwarding-class nc scheduler nc_sch;
        }
    }
}
```

5. Configure the four schedulers (referenced in the scheduler map) that define the four output queues for the scheduler node:

```
[edit]
class-of-service {
    schedulers { # Define scheduling characteristics of each queue
        be_sch { # Transmit rate and buffer management parameters
            transmit-rate percent 10;
            buffer-size remainder;
            priority low;
        }
        ef_sch { # Transmit rate and buffer management parameters
            ...
        }
        af_sch { # Transmit rate and buffer management parameters
            ...
        }
        nc_sch { # Transmit rate and buffer management parameters
            ...
        }
    }
}
```

6. Enable ANCP to communicate with the DSLAM to adjust the CoS shaping rate for the scheduler node.

   You must enable the ANCP feature for performing CoS traffic shaping adjustments, configure the DSLAM as an ANCP neighbor, and specify the DSLAM-assigned identifier for the subscriber local loop represented by the scheduler node:

```
[edit]
```

 Copyright © 2013, Juniper Networks, Inc.

```
protocols {
  ancp {
    qos-adjust; # Enable ANCP to adjust CoS shaping rates
    neighbor 10.2.3.4; # Configure the DSLAM as an ANCP neighbor
    interfaces { # Identify subscribers for which ANCP can adjust shaping rates
      interface-set {
        ifset-of-logical-interfaces {
          access-identifier “dslam port 2/3”; # DSLAM ID for the local loop
        }
      }
    }
  }
}
```

NOTE: If ANCP is not yet enabled, the process starts when you commit a configuration that contains the **protocols ancp** stanza.

7. You can display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical CoS, issue the **show class-of-service interface-set** operational command.

Related Documentation

- Hierarchical CoS Shaping-Rate Adjustments Overview on page 1080
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 1082
- Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 1083
- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1099

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

PART 14

Subscriber Service Activation

- Subscriber Service Activation Overview on page 1121

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 62

Subscriber Service Activation Overview

- CLI-Activated Subscriber Services on page 1121
- Activating and Deactivating Subscriber Services Locally with the CLI on page 1122
- Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers on page 1125
- Default Subscriber Service Overview on page 1127
- Configuring a Default Subscriber Service on page 1127

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

Copyright © 2013, Juniper Networks, Inc.

upgrade qualified subscribers to an advanced service, in addition to retaining the initial 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 552.

Related Documentation

- Activating and Deactivating Subscriber Services Locally with the CLI on page 1122
- Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers on page 1125
- Default Subscriber Service Overview on page 1127

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 552.

Copyright © 2013, 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: larry@isp5.net
Stripped username: larry
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: larry@isp5.net
Stripped username: larry
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
```

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: larry@isp5.net
Stripped username: larry
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: larry@isp5.net
Stripped username: larry
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
```

Related Documentation

- CLI-Activated Subscriber Services on page 1121
- Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers on page 1125
- Default Subscriber Service Overview on page 1127

Copyright © 2013, Juniper Networks, Inc.

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 © 2013, 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 1126
- Using the CLI to Modify a Traffic-Control Profile for a Specific Current Subscriber on page 1126

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 © 2013, 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.

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.

Related Documentation

- Configuring a Default Subscriber Service on page 1127
- CLI-Activated Subscriber Services on page 1121
- Activating and Deactivating Subscriber Services Locally with the CLI on page 1122
- Understanding Gx-Plus Interactions Between the Router and the PCRF on page 543

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 1127
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

Copyright © 2013, Juniper Networks, Inc.

PART 15

Dynamic Firewall Filters, Service Sets and HTTP Redirect for Subscriber Access

- Dynamic Firewall Filters and Service Sets Overview on page 1131
- Parameterized Filters on page 1153
- Configuring Filters for Dynamic Profiles on page 1169
- Configuring Fast Update Filters on page 1183
- Configuring Service Sets in Dynamic Profiles on page 1197
- Firewall Filter Examples on page 1199
- Redirecting HTTP Requests Overview on page 1225
- Configuring HTTP Redirect on page 1229
- HTTP Redirect Examples on page 1233

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 63

Dynamic Firewall Filters and Service Sets Overview

- Dynamic Firewall Filters Overview on page 1132
- Classic Filters Overview on page 1133
- Basic Classic Filter Syntax on page 1135
- Ascend-Data-Filter Policies for Subscriber Management Overview on page 1136
- Ascend-Data-Filter Attribute Fields on page 1138
- Firewall Filters and Enhanced Network Services Mode Overview on page 1141
- Hierarchical Policer Overview on page 1142
- Hierarchical Policer as Filter Action on page 1143
- Interface-Shared Filters Overview on page 1144
- Enhanced Policer Statistics Overview on page 1144
- Fast Update Filters Overview on page 1146
- Basic Fast Update Filter Syntax on page 1149
- Match Conditions and Actions in Fast Update Filters on page 1150
- Dynamic Service Sets Overview on page 1151

Copyright © 2013, Juniper Networks, Inc.

Dynamic Firewall Filters Overview

Firewall filters provide rules that define whether to accept or reject packets that are transiting an interface on a router. The subscriber management feature supports four categories of firewall filters—classic filters, parameterized filters, Ascend-Data-Filters, and fast update filters.

- Classic filters are compiled at commit time and then, when a service is activated, an interface-specific clone of the filter is created and attached to a logical interface. Classic filters are static filters; they cannot contain subscriber-specific terms (also called rules). Classic filters can be applied to interfaces dynamically. This dynamic application is performed by associating input or output filters with a dynamic profile. When triggered, a dynamic profile can apply a named filter or a filter specified in RADIUS to an interface.
- Parameterized filters add the ability to configure firewall filters under a dynamic profile. The filter definitions utilize dynamic-profile variables, which allow you to customize your configuration at session creation time. You can configure a general filter under a dynamic profile and then provide policing rates, destination addresses, ports, and so forth when a dynamic session is activated.
- Ascend-Data-Filters create policies for subscriber traffic. The filter is configured on the RADIUS server and contains rules that specifically match conditions for traffic and define an action for the router to perform.
- Fast update filters are similar to classic filters in many ways. However, fast update filters support subscriber-specific, rather than interface-specific, filter values. Fast update filters also allow individual filter terms to be incrementally added or removed from filters without requiring that the entire filter be recompiled for each modification. Fast update filters are essential for networking environments in which multiple subscribers might share the same logical interface.

You configure firewall filters to determine whether to accept or reject traffic before it enters or exits an interface to which the firewall filter is applied. An *input* (or *ingress*) firewall filter is applied to packets that are entering a network. An *output* (or *egress*) firewall filter is applied to packets that are exiting a network. You can configure firewall filters to subject packets to filtering or class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority).

Related Documentation

- Classic Filters Overview on page 1133
- Ascend-Data-Filter Policies for Subscriber Management Overview on page 1136
- Parameterized Filters Overview on page 1153
- Fast Update Filters Overview on page 1146
- Dynamically Attaching Statically Created Filters for Any Interface Type on page 1170
- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type on page 1169

Copyright © 2013, Juniper Networks, Inc.

- Dynamically Attaching Filters Using RADIUS Variables on page 1172

Classic Filters Overview

The dynamic firewall feature supports classic filters, parameterized filters, and fast update filters. Classic filters are compiled at commit time. When a service activation takes place, the router creates an interface-specific clone of the filter and attaches the clone to the specified logical interface. Classic filters are static filters; they cannot contain subscriber-specific terms, as opposed to fast update filters, which are subscriber-specific. Parameterized filters and policers have their configuration customized at session creation time.

This overview covers:

- Classic Filter Types on page 1133
- Classic Filter Components on page 1133
- Classic Filter Processing on page 1134
- Guidelines for Creating and Applying Classic Filters for Subscriber Interfaces on page 1135

Classic Filter Types

The following classic filter types are supported:

- Port (Layer 2) firewall filter—Port firewall filters apply to Layer 2 switch ports. You can apply port firewall filters only in the ingress direction on a physical port.
- VLAN firewall filter—VLAN firewall filters provide access control for packets that enter a VLAN, are bridged within a VLAN, and leave a VLAN. You can apply VLAN firewall filters in both ingress and egress directions on a VLAN. VLAN firewall filters are applied to all packets that are forwarded to or forwarded from the VLAN.
- Router (Layer 3) firewall filter—You can apply a router firewall filter in both ingress and egress directions on Layer 3 (routed) interfaces.

Classic Filter Components

When creating a classic filter, you first define the family address type (**inet** or **inet6**) and then you define one or more terms that specify the filtering criteria and the action to take when a match occurs.

Each term, or rule, consists of the following components:

- Match conditions—Specifies values or fields that the packet must contain. You can define various match conditions, including:
  - IP source address field
  - IP destination address field
  - Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port field

  - IP protocol field
  - Internet Control Message Protocol (ICMP) packet type
  - TCP flags
  - interfaces

- Actions—Specifies what to do when a match condition occurs. Possible actions are to accept or discard a packet. In addition, packets can be counted to collect statistical information. If no action is specified for a term, the default action is to accept the packet.

Classic Filter Processing

The order of the terms within a classic filter is important. Packets are tested against each term in the order in which the terms are listed in the firewall filter configuration. When a firewall filter contains multiple terms, the router takes a top-down approach and compares a packet against the first term in the firewall filter. If the packet matches the first term, the router executes the action defined by that term to either accept or reject the packet, and no other terms are evaluated. If the router does not find a match between the packet and first term, it then compares the packet to the next term in the firewall filter by using the same match process. If no match occurs between the packet and the second term, the router continues to compare the packet to each successive term defined in the firewall filter until a match is found. If a packet does not match any terms in a firewall filter, the default action is to discard the packet.

You can also specify a precedence (from 0 through 255) for input and output filters within a dynamic profile to force filter processing in a particular order. Setting a lower precedence value for a filter gives it a higher precedence within the dynamic profile. Filters with lower precedence values are applied to interfaces before filters with higher precedence values. A precedence of zero (the default) gives the filter the highest precedence. If no precedence is specified, the filter receives a precedence of zero (highest precedence). Filters with matching precedence (zero or otherwise) are applied in random order.

NOTE: **Dynamic filters do not process outbound packets that are sourced from the routing engine. To filter outbound packets that are sourced from the routing engine, you can create static outbound filters for each interface.**

Copyright © 2013, Juniper Networks, Inc.

Guidelines for Creating and Applying Classic Filters for Subscriber Interfaces

This release supports the dynamic configuration of firewall filters. However, you can also continue to create static firewall filters for interfaces as you do normally, and then dynamically apply those filters to statically created interfaces using dynamic profiles. You can also use dynamic profiles to attach input and output filters through RADIUS.

When creating and applying filters, keep the following in mind:

- This release supports dynamic application of only input and output filters.
- The filters must be interface-specific.
- You can create family-specific **inet** and **inet6** filters.
- You can create interface-specific filters at the **unit** level that apply to any family type (**inet** or **inet6**) configured on the interface.
- You can add or remove both IPv4 and IPv6 filters with the same service activation or deactivation.
- You can remove one filter type without impacting the other type of filter. For example, you can remove IPv6 filters and leave the current IPv4 filters active.
- You can chain up to five input filters and four output filters together.
- If you do not configure and apply a filter, the interface uses the default group filter configuration.
- You cannot modify or delete a firewall filter while subscribers on the same logical interface are bound.

Related Documentation

- Dynamic Firewall Filters Overview on page 1132
- Fast Update Filters Overview on page 1146
- Dynamically Attaching Statically Created Filters for Any Interface Type on page 1170
- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type on page 1169
- Dynamically Attaching Filters Using RADIUS Variables on page 1172
- Verifying and Managing Firewall Filter Configuration on page 1194

Basic Classic Filter Syntax

This section provides the basic classic filter CLI statement syntax. The first part of this syntax provides the CLI statements to associate an input and output filter with a dynamic profile. The second part of this syntax represents the configured input and output filters applied to the dynamic profile. When a DHCP event occurs, the dynamic profile applies the specified filters to the DHCP client interface on the router.

```
[edit]
dynamic-profiles [profile-name] {
  interfaces {
```

```
        [$junos-interface-ifd-name] {
          unit [$junos-underlying-interface-unit] {
            family  family] {
              filter {
                input {
                  [filter-name];
                  precedence [precedence];
                }
                output {
                  [filter-name];
                  precedence [precedence];
                }
              }
            }
          }
        }
      }
    [edit]
    firewall {
      family [family] {
        filter [filter-name] {
          [desired filter configuration]
        }
        filter [filter-name] {
          [desired filter configuration]
        }
      }
    }
```

Related Documentation

- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type on page 1169
- Dynamic Firewall Filters Overview on page 1132

Ascend-Data-Filter Policies for Subscriber Management Overview

Subscriber management enables you to use Ascend-Data-Filters to create policies for subscriber traffic. An Ascend-Data-Filter is a binary value that is configured on the RADIUS server. The filter contains rules that specify match conditions for traffic and an action for the router to perform (such as accept or discard the traffic). The match conditions might include the source and destination IP address or port, the protocol, the filter direction, the traffic class, and policer information.

Subscriber management uses a dynamic profile to obtain the Ascend-Data-Filter attribute (RADIUS attribute 242) from the RADIUS server and apply the policy to a subscriber session. Dynamic profiles support Ascend-Data-Filters for **inet** and **inet6** family types, and both families can be present in a dynamic profile. You include Junos OS predefined variables in the dynamic profiles — **$junos-adf-rule-v4** for family **inet** and **$junos-adf-rule-v6** for **inet6**. The Ascend-Data-Filter attribute can include rules for both address families. The predefined variables map the Ascend-Data-Filter rules for the respective family to the Junos OS firewall filter process. A firewall filter is created and attached to the subscriber's logical interface.

 Copyright © 2013, Juniper Networks, Inc.

You can also configure a static Ascend-Data-Filter by manually entering the required binary data as a hexadecimal string in a dynamic profile. A statically configured Ascend-Data-Filter in a dynamic profile takes precedence over an Ascend-Data-Filter attribute that is received from RADIUS. The static method is time-consuming to configure; it is typically used only for testing purposes.

The Ascend-Data-Filter attribute is supported in RADIUS Access-Accept and Change of Authorization (CoA) messages.

CoA updates existing filters based on the Ascend-Data-Filter Type field, as shown in the following list:

- If the Type field is 1, IPv4 rules are updated and IPv6 rules are unchanged. The opposite is true if the Type field is 3.
- If both Type 1 and 3 are specified, then all rules are updated.
- If the CoA has no Ascend-Data-Filter rules, then the existing rules are unchanged.

Filter Naming Conventions

Each Ascend-Data-Filter has a unique name, which is assigned by the dynamic firewall process, dfwd. The assigned names are displayed in the results of the **show subscriber extensive** and **show firewall** commands. Ascend-Data-Filters use the following naming convention:

__junos_adf_*session#-interfacename-family-direction*

For example:

__junos_adf_33847-ge/1/0/4.53-init-in

Each Ascend-Data-Filter rule maps to a single term, and the term names are simply **t0**, **t1**, ..., **t*n***. If you configure the **counter** option, the router adds a count action to each term that is created. The counter names are a combination of the the term names with **-cnt** appended. For example **t0-cnt** and **t1-cnt**.

Use of Multiple Sessions with Ascend-Data-Filters on an Interface

An interface can have multiple subscriber sessions, each session using its own Ascend-Data-Filter rules. When an Ascend-Data-Filter is applied to a subscriber session, the rules are created independently of any other filters and are added to the interface filter list. The Ascend-Data-Filter rules for the other sessions on the same interface are also added to the filter list. All packets that are processed for the interface must go through all filters, and the filters are applied according to the precedence you set.

Because the filter list can be a combination of several rules, you must consider how the multiple filters coexist. You must ensure that the filters are designed and applied correctly in order to provide the desired filtering and resulting action. For example, a session might have a filter that accepts traffic from Subscriber-A and discards all other traffic. However, a second session on the same interface might have a filter that accepts traffic from Subscriber-B only and discards other traffic. When the two filters are combined in the filter list, traffic from Subscriber-B is discarded by the first filter, and traffic from

Subscriber-A is discarded by the second filter. As a result, no traffic is accepted on the interface because the two filters essentially cancel out each other and discard all traffic.

Optional ADF Filter Requirement for Some Subscribers

When you include either of the predefined variables—**$junos-adf-rule-v4** or **$junos-adf-rule-v6**—in the dynamic profile, by default the RADIUS reply message must include the Ascend-Data-Filter attribute (RADIUS attribute 242) for each subscriber. If the attribute is not included, the router reports an error.

A service provider might apply the same dynamic profile to a mixed pool of subscribers, such that the attribute is included by RADIUS for some of the subscribers and is not included for others. By default, the router returns an error for each of the subscribers without the attribute, consuming system resources. You can configure the dynamic profile to accommodate such a mixture of subscribers by making the attribute requirement optional. To do so, and to suppress attribute error reporting, specify the **not-mandatory** option with the **adf** statement at the **[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter]** hierarchy level. With this configuration, the Ascend-Data-filter is simply not created when the Ascend-Data-Filter attribute is not present.

Related Documentation

- Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions on page 1177
- Ascend-Data-Filter Attribute Fields on page 1138

Ascend-Data-Filter Attribute Fields

Table 103 on page 1138 provides information about the fields used in the Ascend-Data-Filter attribute (RADIUS attribute 242) and how the fields map to Junos OS filter functions. The table lists the fields in the order in which they occur in the Ascend-Data-Filter attribute.

Table 103: Ascend-Data-Filter Attribute Fields

Action or Classifier	Format	Value	Junos OS Filter Function
Type	1 byte	• 1 = IPv4 • 3 = IPv6	
Filter or forward	1 byte	• 0 = filter • 1 = forward	• 0 = maps to **discard** action • 1 = maps to **accept** action
Indirection	1 byte	• 0 = egress • 1 = ingress	• 0 = adds egress terms to the output filter • 1= adds ingress terms to the input filter
Spare	1 byte	–	–

Copyright © 2013, Juniper Networks, Inc.

Table 103: Ascend-Data-Filter Attribute Fields *(continued)*

Action or Classifier	Format	Value	Junos OS Filter Function
Source IP address	IPv4 = 4 bytes IPv6 = 16 bytes	IP address of the source interface	• 0 = no mapping performed • From **source-address** ***address*** entry added to term
Destination IP address	IPv4 = 4 bytes IPv6 = 16 bytes	IP address of the destination interface	• 0 = no mapping performed • From **destination-address** ***address*** entry added to term
Source IP prefix	1 byte	• Type 1 = Number of leading zeros in the wildcard mask • Type 3 = Higher order contiguous bits of the address that make up the network portion of the address	• 0 = no mapping performed • From **source-address** ***prefix*** entry added to term
Destination IP prefix	1 byte	• Type 1 = Number of leading zeros in the wildcard mask • Type 3 = Higher order contiguous bits of the address that make up the network portion of the address	• 0 = no mapping performed • From **destination-address** ***prefix*** entry added to term
Protocol	1 byte	Protocol type	• 0 = no mapping performed • IPv4 = from **protocol** ***number*** added to term • IPv6 = from **next-header** ***number*** added to term
Established	1 byte	Not implemented	Not implemented
Source port	2 bytes	Port number of the source port	From **source-port** ***x - y*** entry added to term
Destination port	2 bytes	Port number of the destination port	From **destination-port** ***x - y*** entry added to term
Source port qualifier	1 byte	• 0 = no compare • 1 = less than • 2 = equal to • 3 = greater than • 4 = not equal to	• 0 = no mapping performed • 1 – 3 = mapped to corresponding option • 4 = mapped to **except** match option

Copyright © 2013, Juniper Networks, Inc.

Table 103: Ascend-Data-Filter Attribute Fields *(continued)*

Action or Classifier	Format	Value	Junos OS Filter Function
Destination port qualifier	1 byte	• 0 = no compare • 1 = less than • 2 = equal to • 3 = greater than • 4 = not equal to	• 0 = no mapping performed • 1 – 3 = mapped to corresponding match option • 4 = mapped to **except** match option
Reserved	2 bytes	Not used	Not used
Marking value	1 byte	• For IPv4 = Type of Service (ToS) • For IPv6 = Differentiated Services Code Point (DSCP)	Not implemented
Marking mask	1 byte	0 = no packet marking	Not implemented
Traffic class	1–41 bytes	• 0 = no traffic class (required if there is no profile) • First byte specifies the length of the ASCII name of the traffic class • Traffic class must be statically configured • Name can optionally be null terminated, which consumes 1 byte • If a name is given, it must match one of the default forwarding classes (such as best-effort) or the name of a forwarding class configured under the **[edit class-of-service scheduler-maps *map-name*]** stanza.	Maps to the forwarding class name. The action **forwarding-class *name*** is added to term.

Copyright © 2013, Juniper Networks, Inc.

Table 103: Ascend-Data-Filter Attribute Fields *(continued)*

Action or Classifier	Format	Value	Junos OS Filter Function
Rate-limit profile	1–41 bytes	• 0 = no rate limit (required if there is no profile) • First byte specifies the length of the ASCII, followed by the ASCII name of the profile • Profile must be statically configured • Name can optionally be null terminated, which consumes 1 byte • If a name is given, it must match the name of one of the firewall policers that is configured under the **[edit firewall]** stanza.	Maps to the **policer** ***policer-name*** action modifier of the same name. The action **policer** ***name*** is added to term.

Related Documentation

- Ascend-Data-Filter Policies for Subscriber Management Overview on page 1136

Firewall Filters and Enhanced Network Services Mode Overview

Under normal conditions, every firewall filter is generated in two different formats -- compiled and term-based. The compiled format is used by the routing engine (RE) kernel, FPCs, and MS-DPs. The term-based format is used by MPCs. Compiled firewall filters are duplicated for each interface or logical interface to which they are applied. Term-based filters, instead of being duplicated, are referenced by each interface or logical interface.

When a combination of MPCs and any other cards populate a chassis, the creation of both firewall filter file formats is necessary. In most networks, the creation of both filter formats and any amount of duplication for compiled firewall filters has no effect on the router. However, in subscriber management networks that include thousands of statically configured subscriber interfaces, creating filters in multiple formats and duplicating those filters for each interface can utilize a large portion of router memory resources. You can use either Enhanced IP Network Services mode or Enhanced Ethernet Network Services mode to improve the scaling and performance specific to routing filters in a subscriber access network that uses statically configured subscriber interfaces.

In configurations where interfaces are created either statically or dynamically and firewall filters are applied dynamically, you must configure the chassis network services to run in enhanced mode. In configurations where interfaces are created statically and firewall filters are applied statically, you must configure chassis network services to run in enhanced mode and also configure each firewall filter for enhanced mode.

Table 104 on page 1142 shows the configuration options when determining enhanced network services mode usage.

Table 104: Enhanced Network Services Mode and Firewall Filter Use Case Determination

Interface and Filter Configuration	Chassis Enhanced Mode Required	Firewall Filter Enhanced Mode Required
Dynamically-created interfaces and dynamically-applied filters	Yes	No
Statically-created interfaces and dynamically-applied filters	Yes	No
Statically-created interfaces and statically-applied filters	Yes	Yes

To achieve significant resource savings for the router, combine chassis and filter enhanced mode configuration as follows:

- Install only MPCs in the chassis.

  NOTE: Configuring chassis network services to run one of the enhanced network services modes results in the router enabling only MPCs and MS-DPCs. Because MS-DPCs use compiled firewall filter format, a router chassis that is configured for one of the enhanced network services modes, configuring standard (non-enhanced) firewall filters for use with any MS-DPCs can decrease optimal resource efficiency.

- When configuring static interfaces on the router, configure chassis network services to run either Enhanced IP Network Services mode or Enhanced Ethernet Network Services mode.
- When statically applying firewall filters to statically-created interfaces, configure any firewall filters for enhanced mode to limit the filter creation to only term-based format.

  NOTE: Any firewall filters that are not configured for enhanced mode are created in both compiled and term-based format, even if the chassis is running one of the enhanced network services modes.

Related Documentation

- *Network Services Mode Overview* in the *Junos OS Administration Library for Routing Devices*
- *Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers* in the *Junos OS Administration Library for Routing Devices*
- Configuring a Filter for Use with Enhanced Network Services Mode on page 1176

Hierarchical Policer Overview

Hierarchial policers rate-limit premium traffic separately from the aggregate traffic on an interface as determined by different configured rates. Hierarchical policing uses two token buckets to maintain two rates: an aggregate and a high priority rate, such as 10Mbps

Copyright © 2013, Juniper Networks, Inc.

and 2Mbps. The traffic is marked differently based on the class of service. Two classes of service are defined for this use: expedited forwarding (EF) and non-expedited forwarding (non-EF). The EF traffic has a user-selectable rate, such as 2Mbps, that is guaranteed before being subject to marking. If there is no EF traffic present, then the non-EF traffic can use up to the 10Mbps rate before being marked. If there is EF traffic present, then the EF traffic is assured up to the 2Mbps (from the 10Mbps) before it becomes subject to marking, but also consumes from the non-EF rate. In this example the EF traffic is guaranteed the 2Mbps and the non-EF traffic has the remaining 8Mbps before being marked.

Hierarchical Policing has the following characteristics:

- Ingress traffic is first classified into premium and non-premium traffic prior to applying a policer.
- The hierarchal policer contains two policers: premium and aggregate. The premium traffic is policed by both the premium policer and aggregate policer. Although premium policer rate-limits the premium traffic, the aggregate policer only decrements the credits but does not drop the packets. The non-premium traffic is rate-limited only by the aggregate policer. Therefore, the premium traffic is assured to have the bandwidth configured for premium and the non-premium traffic is policed to the remaining bandwidth.

Related Documentation

- Hierarchical Policer as Filter Action on page 1143

Hierarchical Policer as Filter Action

Hierarchical policer as filter action enables you to have hierarchial policers as one type of filter action. This is useful in provider edge applications using aggregate policing for general traffic and to apply a separate policer for premium traffic on a logical or physical interface. An interface-specific filter can have a hierarchical policer as a filter action whether or not the hierachical policer is a logical interface policer. A non-interface-specific filter can only have a hierarchical policer without using logical interface-specific as a filter action. The following table summarizes where you can use an interface-specific filter.

Interface-specific Filter	Hierarchical Policer Logical-interface-policer	Allowed
no	no	yes
no	yes	no
yes	no	yes
yes	no	yes

To enable all hierarchical policers of the same name in one filter to share the same policer instance in PFE, use the **filter-specific** statement at the **[edit firewall]** hierarchy level.

Related Documentation

- Hierarchical Policer Overview on page 1142

Interface-Shared Filters Overview

Interface-shared filters can be defined statically or dynamically, but can only be applied using dynamic profiles, and are supported for both client and service sessions. The same interface-shared instance can be attached to multiple interfaces only if these interfaces reference the same interface-shared filter name and have the same shared-name. The shared-name can either be populated from **$junos-interface-set-name**, where the related client session provides the value, or from a service session variable.

With VLAN subscriber interfaces that use the agent-circuit-identifier information, many subscribers share the same underlying logical interface. Because some of these subscribers are related to each other as part of the same household, you must apply an interface-shared filter to the subscriber logical interfaces that make up the household to be able to filter and police these related subscribers at a household level. All interfaces that share the same interface-shared filter instance share the same set of counters and policer actions.

The base filter name of a parameterized filter is assigned depending upon the profile name and the contents of the filter definition. Therefore, when an interface-shared filter is used with parameterized filters, all service sessions that want to share the same instance of an interface-shared filter must have the exact same parameterized filter and profile. A service session uses a different instance of the interface-shared filter if either the parameterized filter or the profile is different.

Related Documentation

- Example: Interface-Shared Filter Configuration on page 1217

Enhanced Policer Statistics Overview

Enhanced policer statistics enable you to display additional statistics for policers as follows:

- Offered packet statistics for traffic subjected to policing.
- OOS packet statistics for packets that are marked out-of-specification (out-of-spec) by the policer. Changes to all packets that have out-of-spec actions, such as discard, color marking, or forwarding-class, are included this counter.
- Transmitted packet statistics for traffic that is not discarded by the policer. When the policer action is discard, the statistics are the same as the in-spec statistics; when the policer action is non-discard (loss-priority or forwarding-class), the statistics are included in this counter.

This feature is supported on MPC/MIC interfaces on MX Series routers and Multi-Rate Ethernet Enhanced Queuing IP Services DPC with SFP and XFP.

To enable this feature, include the **enhanced-policer** statement at the **[edit chassis]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

The new **detail** option has been added to the **show firewall**, **show firewall filter** ***filter-name***, and **show policer** commands.

Related Documentation

- *show policer*

Fast Update Filters Overview

The dynamic firewall feature supports classic filters and fast update filters. Fast update filters support subscriber-specific filter values, as opposed to classic filters, which are interface-specific. Fast update filters allow individual filter terms, or rules, to be added or removed from filters without requiring that you recompile the filter after each modification—terms are added and removed when subscriber services are added and removed.

Using the fast update filters feature involves three distinct operations:

1. Creating the filter—You define fast update filters under the **[edit dynamic-profiles *profile-name* firewall family *family*]** hierarchy. The **dynamic-profiles** stanza enables you to use dynamic variables to create subscriber-specific configurations for the filter's match terms. See "Configuring Fast Update Filters" on page 1183.
2. Associating the filter with a dynamic profile—You use the **[edit dynamic-profiles *profile-name* interface *interface-name* unit *unit-number* family *family*** hierarchy to associate the filter with a dynamic profile. This is the same procedure used for classic filters. See "Associating Fast Update Filters with Interfaces in a Dynamic Profile" on page 1193.
3. Attaching the filter to an interface—When a subscriber logs in, the dynamic profile instantiates the subscriber session and applies the properties of the profile, including the fast update filter, to the session interface. This is the same procedure used for classic filters. Also, similar to classic filters, the name of fast update filters can be provided in a user's RADIUS file.

When a dynamic profile instantiates a subscriber session and applies a fast update filter, the router verifies that the filter is not already present on the session interface. If the filter is not present, the router adds the filter. If the filter is already present on the interface, the router simply adds any new terms that are not in the existing filter. This procedure is reversed when subscriber sessions are deleted. Any terms that were added by a session are then removed when the session is deleted. The filter is deleted when the last subscriber session is deleted.

NOTE: You can optionally specify that a term can be added only once and cannot be modified. See "Match Conditions and Actions in Fast Update Filters" on page 1150.

This overview covers:

- Fast Update Filter Components on page 1147
- Fast Update Filter Processing on page 1147
- Fast Update Filter Names on page 1148
- Guidelines for Creating and Applying Fast Update Filters on page 1148

Copyright © 2013, Juniper Networks, Inc.

Fast Update Filter Components

When creating a fast update filter, you define one or more terms that specify the filtering criteria and the action to take when a match occurs.

Each term consists of the following components:

- Match condition—Specifies values or fields that the packet must contain. You can match a maximum of five fields in a fast update filter. A match condition can contain a single value or range. This differs from classic filters, in which terms can have multiple values. However, you can use additional terms to specify multiple ranges. "Fast Update Filter Match Conditions" on page 1186 lists the supported match conditions for fast update filters. The order in which the terms appear in a fast update filter is not important, because the router examines the most specific term first. (Classic filters examine the terms in the order in which the terms are listed.)
- Action—Specifies what to do when a packet matches the match condition. If no action is specified for a term, the default action is to accept the packet. "Fast Update Filter Actions and Action Modifiers" on page 1187 lists the supported actions for fast update filters.

Terms that are added to the filter during session instantiation must have a unique set of match conditions. Two terms overlap, or conflict, if a packet can match both sets of conditions—as a result, there are two different actions for the packet. You can ensure that terms are unique by using the **$junos-subscriber-ip-address** variable as the **source-address** (for an input filter) or **destination-address** (for an output filter) in the **from** statement. You must then supply the **source-address** or **destination-address** condition, as appropriate, as the first condition in the **match-order** statement.

Related Documentation

- Fast Update Filter Actions and Action Modifiers on page 1187
- Fast Update Filter Match Conditions on page 1186
- Avoiding Conflicts When Terms Match on page 1188

Fast Update Filter Processing

You must use the **match-order** statement to explicitly specify the order in which the router examines filter match conditions. Also, the router examines only those conditions that you include in the **match-order** statement. When a fast update filter contains multiple terms, the router compares a packet against the terms starting with the most specific condition first. When the packet first matches a condition, the router performs the action defined in the term to either accept or reject the packet, and then no other terms are evaluated. If the router does not find a match between the packet and first term, it then compares the packet to the next term in the filter. The router continues to compare the packet to the next specified term until a match is found. If there is no match after all terms have been examined, the router silently drops the packet.

You can specify a precedence (from 0 through 255) for input and output filters within a dynamic profile to force filter processing in a particular order. Setting a lower precedence value for a filter gives it a higher precedence within the dynamic profile. Filters with lower

Copyright © 2013, Juniper Networks, Inc.

precedence values are applied to interfaces before filters with higher precedence values. A precedence of zero (the default) gives the filter the highest precedence. If no precedence is specified, the filter receives a precedence of zero (highest precedence). Filters with matching precedence (zero or otherwise) are applied in random order.

Fast Update Filter Names

When a filter is attached to an interface, the router first searches for a classic filter with the specified name, and then uses the classic filter. If no classic filter exists with that name, the router then searches in the dynamic profile for a fast update filter with the specified name, and uses that filter.

If two different dynamic profiles include a fast update filter with the same name, the **match-order** specification of the two filters must be identical. If the two filters are activated on the same interface, the terms are added together.

The router includes the filter name in **show firewall** command results. The router also creates unique names for filter terms and counters for the **show firewall** command.

When a fast update filter is created by the activation of a dynamic profile, the router creates an interface-specific name for the filter. The name uses the following format, which is also used for classic filters:

<filter-name>-<interface-name>.<subunit>-<direction>

For example, an input filter named **httpFilter** on interface **ge-1/0/0.5** is named as follows (**in** indicates an input filter and **out** indicates an output filter):

http-filter-ge-1/0/0.5–in

The router creates unique names for the filter terms and counters by appending the session ID to all term and counter names. Terms that use the **only-at-create** statement have a session-id of 0. Terms and counters use the following format:

<term-name>-<session-id>

<counter-name>-<session-id>

Guidelines for Creating and Applying Fast Update Filters

Fast update filters enable you to create subscriber-specific firewall filters and dynamically apply these filters to statically created interfaces using dynamic profiles. Individual terms can be added to, or removed from, a filter without requiring that the entire filter be recompiled.

When creating and applying fast update filters, keep the following in mind:

- This release supports dynamic application of input and output filters.
- You cannot use the same fast update filter as both an input and output filter in the same dynamic profile attached to an interface.
- Fast update filters must always include terms that permit DHCP traffic to pass. See "Configuring Filters to Permit Expected Traffic" on page 1187.

 Copyright © 2013, Juniper Networks, Inc.

- You can create **family inet** and **inet6** filters.
- You can add or remove both IPv4 and IPv6 filters with the same service activation or deactivation.
- You can remove one filter type without impacting the other type of filter. For example, you can remove IPv6 filters and leave the current IPv4 filters active.
- The **interface-specific** statement is required for all fast update filters.
- The **match-order** statement is required—you must explicitly state the order of the match fields in a fast update filter. See "Configuring the Match Order for Fast Update Filters" on page 1184.
- The **match-order** statement uses an implied wildcard for conditions that you specify in the statement. If you specify a condition that is not also configured in the **from** specification of a filter term, the router considers that a wildcard for that condition.
- A filter term can have only a single value or range; however, you can configure multiple terms to specify multiple ranges.
- You can match a maximum of five match conditions in a filter.

Related Documentation

- Dynamic Firewall Filters Overview on page 1132
- Classic Filters Overview on page 1133
- Dynamically Attaching Statically Created Filters for Any Interface Type on page 1170
- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type on page 1169
- Verifying and Managing Firewall Filter Configuration on page 1194

Basic Fast Update Filter Syntax

This section shows the basic fast update filter statement syntax. The first part of this syntax provides the CLI statements to associate an input and output filter with a dynamic profile. The second part of this syntax represents the configured input and output filters associated to the dynamic profile. When a DHCP event occurs, the dynamic profile applies the specified filters to the DHCP client interface on the router.

```
[edit dynamic-profiles profile-name]
interfaces {
  $junos-interface-ifd-name {
    unit $junos-underlying-interface-unit {
      family family {
        filter {
          input filter-name;
          precedence precedence;
          output filter-name;
          precedence precedence;
        }
      }
    }
  }
```

```
}
[edit dynamic-profiles profile-name]
firewall {
  family family {
    fast-update-filter filter-name {
      [desired filter configuration]
    }
    fast-update-filter filter-name {
      [desired filter configuration]
    }
  }
}
```

Related Documentation

- Configuring Fast Update Filters on page 1183

Match Conditions and Actions in Fast Update Filters

To create a fast update filter, you use the **term** statement to specify conditions that a packet must have, and to specify the action the router performs when those conditions exist in the packet.

This section covers:

- Match Conditions on page 1150
- Actions on page 1151
- Adding Terms Only Once on page 1151

Match Conditions

Match conditions specify characteristics that a packet must have—if the conditions exist in the packet, the router then performs the specified action. You use the **from** keyword in the **term** statement to specify match conditions for the filter. The packet must match all conditions in the **from** specification for the action to be performed, which also means that their order in the **from** specification is not important.

An individual condition in a **from** specification can contain a single value or range. You can match a maximum of five match conditions in a filter.

"Fast Update Filter Match Conditions" on page 1186 lists the match conditions you can use in fast update filters.

NOTE: The router uses an implied wildcard for conditions that you include in the **match-order** statement. If you include a condition that is *not* configured in the **from** specification of a filter term, the router considers that a wildcard for the condition.

For example, if you include the **dscp** condition in the **match-order** statement, but do not configure a **dscp** value in the **from** specification of the filter term, the router performs the action configured in the **then** specification of the filter on all DSCP values.

Copyright © 2013, Juniper Networks, Inc.

Actions

Actions and action modifiers specify the operation the router performs when a particular match condition exists in a packet. You use the **then** keyword in the **term** statement to specify the actions to perform on packets whose characteristics match the conditions specified in the preceding **from** specification.

Action modifiers are actions taken in addition to the specified action. You can configure any combination of action modifiers. For the action or action modifier to take effect, all conditions in the **from** specification must match. If you specify **log** as one of the actions in a term, this constitutes a termination action; whether any additional terms in the filter are processed depends on the traffic through the filter. The action modifier operations carry a default **accept** action. For example, if you specify an action modifier and do not specify an action, the specified action modifier is implemented and the packet is accepted.

"Fast Update Filter Actions and Action Modifiers" on page 1187 lists the actions and action modifiers you can use in fast update filters.

Adding Terms Only Once

You can optionally specify that a term can be added only when the fast update filter is first created, and cannot be later changed by adding or removing conditions. We recommend that you only use the **only-at-create** option for terms that do not include subscriber-specific data in their match conditions, such as common or default terms (counting the default drop packet, for instance).

Related Documentation

- Configuring Terms for Fast Update Filters on page 1185
- Fast Update Filter Match Conditions on page 1186
- Fast Update Filter Actions and Action Modifiers on page 1187

Dynamic Service Sets Overview

A service set is a collection of services to be performed by an Adaptive Services (AS) or Multiservices PIC. You configure a service-set definition at the **[edit services]** hierarchy level. You can then apply the service set to one or more interfaces on the router. The service set can be applied either dynamically or statically.

To dynamically associate a service set to interfaces you include the **service-set** statement with the **input** or **output** statement at the **[edit dynamic-profiles** ***profile-name*** **interfaces** ***interface-name*** **unit** ***logical-unit-number*** **family** ***family*** **service]** hierarchy level.

To statically associate a defined service set with an interface, you include the **service-set** statement with the **input** or **output** statement at the **[edit interfaces** ***interface-name*** **unit** ***logical-unit-number*** **family** ***family*** **service]** hierarchy level.

Related Documentation

- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197
- Verifying and Managing Service Sets Information on page 1198

Copyright © 2013, Juniper Networks, Inc.

- For information about creating service sets, see “Service Set Configuration Guidelines” in the *Junos OS Services Interfaces Library for Routing Devices*.
- For information about statically applying service sets to interfaces, see “Applying Filters and Services to Interfaces” in the *Junos OS Services Interfaces Library for Routing Devices*.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 64

Parameterized Filters

- Parameterized Filters Overview on page 1153
- Basic Parameterized Filter Syntax on page 1154
- Unique Identifiers for Firewall Variables in Dynamic Profiles on page 1154
- Sample Dynamic-Profile Configuration for Parameterized Filters on page 1156
- Dynamic Profile Configuration and UID Substitution Comparison for Parameterized Filters on page 1158
- Example: Dynamic-Profile Parsing on page 1162
- Parameterized Filters Configuration Considerations on page 1163
- Guidelines for Creating and Applying Parameterized Filters for Subscriber Interfaces on page 1164
- Parameterized Filter Processing Overview on page 1165
- IPv4 Parameterized Filter Match Conditions on page 1166
- IPv6 Parameterized Filter Match Conditions on page 1167
- Parameterized Filter Actions and Modifiers on page 1168
- Parameterized Filter Policer Actions on page 1168

Parameterized Filters Overview

Parameterized filters are configured under a dynamic profile. The filter definition can utilize dynamic-profile variables, allowing their configuration to be customized at session creation time. The user can configure a general baseline filter under a dynamic profile and then provide specific variables of that filter when a dynamic session is activated. These variables can include policing rates, destination addresses, ports, and other items.

In order to provide better scaling, the system analyzes a dynamic profile, and then determines whether the set of variables for one session is the same as for a previous session. Each set of variables is assigned a unique identifier (UID). If a matching filter already exists, the session creates an interface-specific filter copy of that filter template. If the filter does not already exist, the session reads the configuration and compiles a new filter. This filter is installed as a template with an interface-specific filter copy for the current session pointing at it.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Dynamic Firewall Filters Overview on page 1132
- Verifying and Managing Firewall Filter Configuration on page 1194
- Unique Identifiers for Firewall Variables in Dynamic Profiles on page 1154
- Sample Dynamic-Profile Configuration for Parameterized Filters on page 1156
- *Dynamic Profile After UID Substitutions for Parameterized Filters*
- Dynamic Profile Configuration and UID Substitution Comparison for Parameterized Filters on page 1158
- Example: Dynamic-Profile Parsing on page 1162
- Parameterized Filters Configuration Considerations on page 1163
- Parameterized Filter Processing Overview on page 1165

Basic Parameterized Filter Syntax

Parameterized filter syntax follows the standard Junos OS filter syntax. When a match condition is met, an action is applied.

Related Documentation

- Basic Classic Filter Syntax on page 1135
- Sample Dynamic-Profile Configuration for Parameterized Filters on page 1156

Unique Identifiers for Firewall Variables in Dynamic Profiles

The system uses unique identifiers (UIDs) to aid with scaling. The UID enables the system to determine when configuration objects from multiple subscribers are identical and can be shared. In many situations, such as a filter definition, sharing a single filter program among multiple subscribers instead of creating a new program for every subscriber helps to conserve system resources.

Within a dynamic profile a UID is used to name a configuration object. The system assigns the value of the UID (the object's name) based upon all the variables contained within that configuration stanza along with the dynamic profile's name. The assigned UID value consists of the UID name combined with the string **_UID** and a unique number. For instance, the UID **$my-filter** might be given the value **my-filter_UID1022**.

You must first define a UID under the **variable** stanza using the option **uid**. The UID must be defined at the end, after all the variables that are assigned values externally.

```
dynamic-profile test-profile {
  [variables] {
    ... [other variables] ...
    [my-filter] {
      uid;
    }
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

After a UID has been defined, it can then be used to name an object:

```
dynamic-profile test-profile {
  firewall {
    family inet {
      filter [$my-filter] {
        ... [filter definition that makes use of other variables] ...
      }
    }
  }
}
```

As previously described, the system assigns the value of **$my-filter** depending on the values of the variables used within that filter's definition.

The UID is also used in any other place that the object's name is used. For example, here is an interface stanza to use **$my-filter** as an input filter:

```
dynamic-profile [test-profile] {
  interfaces {
    [$junos-interface-ifd-name]" {
      unit [$junos-interface-unit] {
        family inet {
          filter {
            input [$my-filter];
          }
        }
      }
    }
  }
}
```

You can define multiple configuration objects of the same type (that is, multiple filters) as long as each one uses its own, individual, UID. To ensure that the system selects the correct object when assigning a name, use the **uid-reference** variable.

When the uid-reference is used, it is effectively evaluated twice. First, the value of the uid-reference variable is retrieved. Second, that value is used as the name of a UID and that UID value is retrieved. A uid-reference with a value that is not the name of a UID is considered an error.

A uid-reference is defined similarly to any other variable:

```
dynamic-profile [test-profile] {
  variables {
    [my-filter-selector] {
      uid-reference;
    }
  }
}
```

A uid-reference is used wherever the name of the object is needed. One example is the name of the input filter in the following interface stanza:

```
dynamic-profile [test-profile] {
  interfaces {
```

```
        [$junos-interface-ifd-name] {
          unit [$junos-interface-unit] {
            family inet {
              filter {
                input [$my-filter-selector];
              }
            }
          }
        }
      }
    }
```

Consider the case where two parameterized filters are defined: **$my-filter-1** and **$my-filter-2**. The **$my-filter-selector** variable might be assigned the value **my-filter-1** or **my-filter-2**, depending upon which filter is appropriate.

Related Documentation

- Parameterized Filter Processing Overview on page 1165
- Parameterized Filters Configuration Considerations on page 1163

Sample Dynamic-Profile Configuration for Parameterized Filters

In the following sample configuration, the **my-svc-prof** profile provides two different filters: **my-filt-1gw** and **my-filt-2gw**. These filters match on either one or two gateway addresses and apply a policer for that traffic. The name of the filter to apply, the gateway addresses, and the bandwidth for the policer are passed into the service profile from the RADIUS service activation. The uid-reference type supports selection of a particular UID generated object out of multiple objects in the profile. The UID type indicates that a variable is used for UID generation.

```
dynamic-profile {
  [my-svc-prof] {
    variable {
      [my-in-filter] {
        mandatory;
        uid-reference;
      }
      gw1 {
        mandatory;
      }
      gw2 {
        mandatory;
      }
      bw {
        mandatory;
      }
      my-filt-1gw {
        uid;
      }
      my-filt-2gw {
        uid;
      }
      [my-policer] {
        uid;
```

 Copyright © 2013, Juniper Networks, Inc.

```
    }
}
interfaces {
    [$junos-interface-ifd-name] {
        unit [$junos-underlying-interface-unit] {
            family inet {
                filter {
                    input [$my-in-filter];
                }
            }
        }
    }
}
firewall {
    policer [$my-policer] {
        if-exceeding {
            bandwidth-limit $bw;
            burst-size-limit 15000;
        }
        then discard;
    }
    family inet {
        filter [$my-filt-1gw] {
            interface-specific;
            term t0 {
                from {
                    destination-address $gw1;
                }
                then {
                    policer [$my-policer];
                }
            }
            term last {
                then {
                    count drops;
                    discard;
                }
            }
        }
        filter [$my-file-2gw] {
            interface-specific;
            term t0 {
                from {
                    destination-address {
                        $gw1;
                        $gw2;
                    }
                }
                then {
                    policer [$my-policer];
                }
            }
            term last {
                then {
                    count drops;
                    discard;
```

```
                    }
                  }
                }
              }
            }
          }
        }
```

Related Documentation

- *Dynamic Profile After UID Substitutions for Parameterized Filters*
- Dynamic Profile Configuration and UID Substitution Comparison for Parameterized Filters on page 1158
- Example: Dynamic-Profile Parsing on page 1162

Dynamic Profile Configuration and UID Substitution Comparison for Parameterized Filters

Table 105 on page 1158 compares "Sample Dynamic-Profile Configuration for Parameterized Filters" on page 1156 and *Dynamic Profile After UID Substitutions for Parameterized Filters* when the following service is activated.

```
my-svc-prof(my-filt-1gw, 207.17.137.239/32, 0, 5m)
```

Table 105: Dynamic Profiles and UID Substitution Comparison

Dynamic Profile Configuration	Result After Substitution	Comment
dynamic-profile {	dynamic-profile {	-
[*my-svc-prof*] {	[*my-svc-prof*] {	-
variable {	-	Define the variables.
[*my-in-filter*] {	-	-
mandatory;	-	-
uid-reference;	-	Assign the name of a UID variable to *my-in-filter.*.
}	-	-
gw1 {	-	-
mandatory;	-	-
}	-	-
gw2 {	-	-
mandatory;	-	-

Copyright © 2013, Juniper Networks, Inc.

Table 105: Dynamic Profiles and UID Substitution Comparison *(continued)*

Dynamic Profile Configuration	Result After Substitution	Comment
}	-	-
bw {	-	-
mandatory;	-	-
}	-	-
[*my-filt-1gw*] {	-	-
uid;	-	Type is a UID.
}	-	-
[*my-filt-2gw*] {	-	-
uid;	-	Type is a UID.
}	-	-
[*my-policer*] {	-	-
uid;	-	Type is a UID.
}	-	-
}	-	-
	-	-
interfaces {	interfaces {	-
[*$junos-interface-ifd-name*] {	ge-1/0/0 {	-
unit [*$junos-underlying-interface-unit*] {	unit 7 {	-
family inet {	family inet {	-
filter {	filter {	-
input [*$my-in-filter*];	input my-filt-1gw_UID1022;	Substitute the value of my-filt-1gw for *my-in-filter,* but because my-filt-1gw is a UID reference, substitute the value of *$my-filt-1gw*: my-filt-1gw_UID1022.
}	}	-

Table 105: Dynamic Profiles and UID Substitution Comparison *(continued)*

Dynamic Profile Configuration	Result After Substitution	Comment
}	}	-
}	}	-
}	}	-
}	}	-
		-
firewall {	firewall {	-
policer [*$my-policer*] {	policer my-policer_UID1005 {	Substitute UID name.
if-exceeding {	if-exceeding {	-
bandwidth-limit $bw;	bandwidth-limit 5m;	-
burst-size-limit 15000;	burst-size-limit 15000;	-
}	}	-
then discard;	then discard;	-
}	}	-
family inet {	family inet {	-
filter [*$my-filt-1gw*] {	filter my-filt-1gw_UID1022 {	Substitute UID name
interface-specific;	interface-specific;	-
term t0 {	term t0 {	-
from {	from {	-
destination-address $gw1;	destination-address 207.17.137.239/32;	Substitute $gw1 value.
}	}	-
then {	then {	-
policer [*$my-policer*];	policer my-policer_UID1005;	Substitute UID name.
}	}	-
}	}	-

Copyright © 2013, Juniper Networks, Inc.

Table 105: Dynamic Profiles and UID Substitution Comparison *(continued)*

Dynamic Profile Configuration	Result After Substitution	Comment
term last {	term last {	-
then {	then {	-
count drops;	count drops;	-
discard;	discard;	-
}	}	-
}	}	-
}	}	-
filter [*$my-file-2gw*] {	filter my-filt-2gw_UID1018 {	Substitute UID name
interface-specific;	interface-specific;	-
term t0 {	term t0 {	-
from {	from {	-
destination-address {	destination-address {	-
$gw1;	207.17.137.239/32;	Substitute $gw1 value
$gw2;	0;	Substitute $gw2 value
}	}	-
}	}	-
then {	then {	-
policer [*$my-policer*];	policer my-policer_UID1005;	Substitute UID name
}	}	-
}	}	-
term last {	term last {	-
then {	then {	-
count drops;	count drops;	-
discard;	discard;	-

Table 105: Dynamic Profiles and UID Substitution Comparison *(continued)*

Dynamic Profile Configuration	Result After Substitution	Comment
}	}	-
}	}	-
}	}	-
}	}	-
}	}	-
}	}	-
}	}	-

Related Documentation

- Sample Dynamic-Profile Configuration for Parameterized Filters on page 1156
- *Dynamic Profile After UID Substitutions for Parameterized Filters*
- Example: Dynamic-Profile Parsing on page 1162

Example: Dynamic-Profile Parsing

The following example shows the basic dynamic-profile parsing steps for "Dynamic Profile Configuration and UID Substitution Comparison for Parameterized Filters" on page 1158. The steps apply to any parameterized filter.

1. Read **dynamic-profiles my-svc-prof interface ge-1/0/0 unit 7 family inet filter input** and get the value **my-filt-1gw_UID1022**. The **my-in-filter** variable received the name of the UID (**my-filt-1gw**) from the first service parameter. The name **my-filt-1gw_UID1022** comes from the value of the **my-filt-1gw UID**.
2. Determine whether a static filter called **my-filt-1gw_UID1022** exists. If so, this is the existing classic filter case and not a parameterized filter.
3. Try to read **dynamic-profile my-svc-prof firewall family inet fast-update-filter my-filt-1gw_UID1022**'. If this exists, this is a fast update filter, not a parameterized filter.
4. Try to read **dynamic-profile my-svc-prof firewall family inet filter my-filt-1gw_UID1022**. If this does not exist, return a "filter not found" error.
5. Search for a template named **my-filt-1gw_UID1022**. If it does not exist:
   a. Read the parameterized filter configuration. This adds the match destination address **207.17.137.239** and the policer **my-policer_UID1005** as the action.
   b. Determine whether **my-policer_UID1005** exists. If it does not, read the **dynamic-profile my-svc-prof firewall policer my-policer_UID1005** configuration and create the **my-policer_UID1005** policer.

 Copyright © 2013, Juniper Networks, Inc.

c. Compile the **my-filt-1gw_UID1022** filter.

d. Install **my-filt-1gw_UID1022** as a filter template.

6. Create and install an interface-specific filter reference named my-filt-1gw_UID1022-ge-1/0/0.7-in with **my-filt-1gw_UID1022** as its template.

7. Attach **my-filt-1gw_UID1022-ge-1/0/0.7-in** to interface **ge-1/0/0.7**.

When subsequent sessions are created with the same parameters, the system returns the same **my-filt-1gw_UID1022** filter name. In this case, Step 5 finds the existing filter template and proceeds directly to Step 6.

Related Documentation

- Sample Dynamic-Profile Configuration for Parameterized Filters on page 1156
- *Dynamic Profile After UID Substitutions for Parameterized Filters*
- Dynamic Profile Configuration and UID Substitution Comparison for Parameterized Filters on page 1158

Parameterized Filters Configuration Considerations

Keep the following considerations in mind when configuring parameterized filters.

- Subscriber IP Address on page 1163
- Interaction with Static Configuration on page 1163
- Interface-Specific on page 1164
- Service Session Support on page 1164
- Filter Naming Conventions on page 1164

Subscriber IP Address

In most deployment scenarios, the interface is based on the subscriber's IP address. Because subscribers may not be unique, they cannot be used in determining similar filters and policers. Do not use the **junos-subscriber-ip-address IP** address as a match candidate. Doing so causes unique filters per subscriber, which inhibits scaling.

Interaction with Static Configuration

Searching for a filter to attach takes place in the following order:

1. Static filter. For example, **firewall family inet filter my-filter**.

2. Fast update filter within the current dynamic profile. For example, **dynamic-profile [*profile-name*] firewall family inet fast-update-filter my-filter**.

3. Parameterized filter within the current dynamic profile. For example, **dynamic-profile [*profile-name*] firewall family inet filter**.

The following static configuration objects may be referenced by a parameterized filter. The search order is first in the static configuration and then in the current dynamic-profile:

- firewall policer
- firewall hierarchical-policer
- three-color policer
- policy-options prefix-list

If an object in the static configuration is being used by an active parameterized filter, you cannot delete that object from the configuration while the subscriber is logged in.

Interface-Specific

All dynamic service filters must be defined as interface-specific.

Service Session Support

Parameterized filters and policers are supported for service activations only, not client sessions.

Filter Naming Conventions

The base filter name is based on the interface and direction (ingress and egress) appended to it. With parameterized filters, the filter-naming process comes from the UID.

Related Documentation

- Dynamic Firewall Filters Overview on page 1132
- Verifying and Managing Firewall Filter Configuration on page 1194
- Unique Identifiers for Firewall Variables in Dynamic Profiles on page 1154
- Sample Dynamic-Profile Configuration for Parameterized Filters on page 1156
- *Dynamic Profile After UID Substitutions for Parameterized Filters*
- Dynamic Profile Configuration and UID Substitution Comparison for Parameterized Filters on page 1158
- Example: Dynamic-Profile Parsing on page 1162
- Parameterized Filter Processing Overview on page 1165

Guidelines for Creating and Applying Parameterized Filters for Subscriber Interfaces

This release supports the dynamic configuration of firewall filters. However, you can also continue to create static firewall filters for interfaces as you do normally, and then dynamically apply those filters to statically created interfaces using dynamic profiles. You can also use dynamic profiles to attach input and output filters through RADIUS.

When creating and applying filters, keep the following in mind:

- This release supports dynamic application of only input and output filters.
- The filters must be interface-specific.
- You can create family-specific **inet** and**inet6** filters.

 Copyright © 2013, Juniper Networks, Inc.

- You can create interface-specific filters at the **unit** level that apply to any family type (**inet** or **inet6**) configured on the interface.
- You can add or remove both IPv4 and IPv6 filters with the same service activation or deactivation.
- You can remove one filter type without impacting the other type of filter. For example, you can remove IPv6 filters and leave the current IPv4 filters active.
- You can chain up to five input filters and four output filters together.
- If you do not configure and apply a filter, the interface uses the default group filter configuration.
- You cannot modify a firewall filter while subscribers on the same logical interface are bound.

Related Documentation

- Parameterized Filter Processing Overview on page 1165
- Parameterized Filters Configuration Considerations on page 1163

Parameterized Filter Processing Overview

When creating a parameterized filter, you first define the family address type (**inet** or **inet6**) and then you define one or more terms that specify the filtering criteria and the action to take when a match occurs.

Each term, or rule, consists of the following components:

- Match conditions—Specifies values or fields that the packet must contain. You can define various match conditions, including:
  - IP source address field
  - IP destination address field
  - Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port field
  - IP protocol field
  - Internet Control Message Protocol (ICMP) packet type
  - TCP flags
  - interfaces
- Actions—Specifies what to do when a match condition occurs. Possible actions are to accept or discard a packet. In addition, packets can be counted to collect statistical information. If no action is specified for a term, the default action is to accept the packet.

The processing of parameterized filters is the same as classic filters. The order of the terms within a parameterized filter is important. Packets are tested against each term in the order in which the terms are listed in the firewall filter configuration. When a firewall filter contains multiple terms, the router takes a top-down approach and compares a

Copyright © 2013, Juniper Networks, Inc.

packet against the first term in the firewall filter. If the packet matches the first term, the router executes the action defined by that term to either accept or reject the packet, and no other terms are evaluated. If the router does not find a match between the packet and first term, it then compares the packet to the next term in the firewall filter by using the same match process. If no match occurs between the packet and the second term, the router continues to compare the packet to each successive term defined in the firewall filter until a match is found. If a packet does not match any terms in a firewall filter, the default action is to discard the packet.

You can also specify a precedence (from 0 through 255) for input and output filters within a dynamic profile to force filter processing in a particular order. Setting a lower precedence value for a filter gives it a higher precedence within the dynamic profile. Filters with lower precedence values are applied to interfaces before filters with higher precedence values. A precedence of zero (the default) gives the filter the highest precedence. If no precedence is specified, the filter receives a precedence of zero (highest precedence). Filters with matching precedence (zero or otherwise) are applied in an unspecified order.

NOTE: **Parameterized filters do not support outbound packets that are sourced from the routing engine.**

Related Documentation

- Parameterized Filters Configuration Considerations on page 1163

IPv4 Parameterized Filter Match Conditions

The following IPv4 match conditions are supported for parameterized filters. Their syntax is the same as the static filter syntax.

address
destination-address
destination-port
destination-port-except
destination-prefix-list
dscp
dscp-except
forwarding-class
forwarding-class-except
icmp-code
icmp-code-except
icmp-type
icmp-type-except
loss-priority
loss-priority-except
packet-length
packet-length-except
port
port-except

Copyright © 2013, Juniper Networks, Inc.

precedence
precedence-except
prefix-list
protocol
protocol-except
service-filter-hit
source-address
source-class
source-port
source-port-except
source-prefix-list
ttl
ttl-except

Related Documentation

- *Firewall Filter Match Conditions for IPv4 Traffic*

IPv6 Parameterized Filter Match Conditions

The following IPv6 match conditions are supported for parameterized filters. Their syntax is the same as the static filter syntax.

address
destination-address
destination-port
destination-port-except
destination-prefix-list
forwarding-class
forwarding-class-except
icmp-code
icmp-code-except
icmp-type
icmp-type-except
loss-priority
loss-priority-except
packet-length
packet-length-except
port
port-except
prefix-list
service-filter-hit
source-address
source-class
source-port
source-port-except
source-prefix-list
traffic-class
traffic-class-except

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *Firewall Filter Match Conditions for IPv6 Traffic*

Parameterized Filter Actions and Modifiers

The following actions and modifiers are supported for parameterized filters. Their syntax is the same as the static filter syntax.

accept
count
discard
forwarding-class
hierarchical-policer
log
loss-priority
next
policer
port-mirror
port-mirror-instance
reject
routing-instance
sample
service-accounting
service-accounting-deferred
service-filter-hit
three-color-policer

Related Documentation

- *Firewall Filter Terminating Actions*
- *Firewall Filter Nonterminating Actions*

Parameterized Filter Policer Actions

The following policer actions are supported for parameterized filters. Their syntax is the same as the existing static policer syntax.

discard
forwarding-class
loss-priority

Related Documentation

- *Firewall Filter Terminating Actions*
- *Firewall Filter Nonterminating Actions*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 65

Configuring Filters for Dynamic Profiles

- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type on page 1169
- Dynamically Attaching Statically Created Filters for Any Interface Type on page 1170
- Dynamically Attaching Filters Using RADIUS Variables on page 1172
- Defining Dynamic Filter Processing Order on page 1173
- Configuring Firewall Filter Bypass on page 1174
- Configuring Service Packet Counting on page 1175
- Configuring a Filter for Use with Enhanced Network Services Mode on page 1176
- Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions on page 1177
- Verifying and Managing Dynamic Ascend-Data-Filter Policy Configuration on page 1179
- Configuring Unicast RPF and Fail Filters in Dynamic Profiles for Subscriber Interfaces on page 1179
- Configuring Unicast RPF in Dynamic Profiles for Subscriber Interfaces on page 1180
- Configuring a Fail Filter for Unicast RPF in Dynamic Profiles for Subscriber Interfaces on page 1180

Dynamically Attaching Statically Created Filters for a Specific Interface Family Type

You can dynamically attach statically created filters for either IPv4 (**inet**) or IPv6 (**inet6**) interface types. These filters apply only to interfaces of the specified type.

Before you can attach a statically created filter using a dynamic profile.

1. Create the filters you want to attach.

   See the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices* for detailed information about classic firewall filters and how to create them. See "Configuring Fast Update Filters" on page 1183 for information about creating fast update filters.

2. Create a basic dynamic profile.

   See "Configuring a Basic Dynamic Profile" on page 673.

To dynamically attach statically created input and output filters:

1. Specify the unit family type you want to use when dynamically attaching the filters.

   a. For IPv4 interfaces, specify the **inet** unit family.

      [edit dynamic-profiles access-profile interfaces ge-1/1/1 unit 1]
      user@host# set family inet

   b. For IPv6 interfaces, specify the **inet6** unit family.

      [edit dynamic-profiles access-profile interfaces ge-1/1/1 unit 1]
      user@host# set family inet6

2. Specify the input filter in the dynamic profile.

   [edit dynamic-profiles access-profile interfaces ge-1/1/1 unit 1 family inet]
   user@host# set filter input static-input-filter

3. Specify the output filter in the dynamic profile.

   NOTE: **The following example specifies an optional precedence value for the output filter.**

   [edit dynamic-profiles access-profile interfaces ge-1/1/1 unit 1 family inet]
   user@host# set filter output static-output-filter precedence 50

Related Documentation

- Classic Filters Overview on page 1133
- Fast Update Filters Overview on page 1146
- Dynamically Attaching Statically Created Filters for Any Interface Type on page 1170
- Dynamically Attaching Filters Using RADIUS Variables on page 1172
- For information about Junos OS default groups, see the *CLI User Guide*
- For information about firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*

Dynamically Attaching Statically Created Filters for Any Interface Type

You can dynamically attach statically created filters for any interface type. These filters apply to any interfaces that are created using the dynamic profile.

NOTE: **For an L2TP LNS on MX Series routers, you can attach firewall for static LNS sessions by configuring these at logical interfaces directly on the inline services device (si-fpc/pic/port). RADIUS-configured firewall attachments are not supported.**

Before you can attach a statically created filter using a dynamic profile.

1. Create the filters you want to attach.

 Copyright © 2013, Juniper Networks, Inc.

See the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices* for detailed information about classic firewall filters and how to create them. See "Configuring Fast Update Filters" on page 1183 for information about creating fast update filters.

2. Create a basic dynamic profile.

   See "Configuring a Basic Dynamic Profile" on page 673.

To dynamically attach statically created input and output filters for all interfaces created dynamically using the dynamic profile:

1. Access the dynamic profile, interface, and unit that you want to use when applying the static filters.

   ```
   [edit]
   user@host# edit dynamic-profiles access-profile interfaces ge-1/1/1 unit 1
   ```

2. Specify the input filter for the interface unit.

   ```
   [edit dynamic-profiles access-profile interfaces ge-1/1/1 unit 1]
   user@host# set filter input static-input-filter
   ```

3. Specify the output filter for the interface unit.

   ```
   [edit dynamic-profiles access-profile interfaces ge-1/1/1 unit 1]
   user@host# set filter output static-output-filter
   ```

Related Documentation

- Classic Filters Overview on page 1133
- Fast Update Filters Overview on page 1146
- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type on page 1169
- Dynamically Attaching Filters Using RADIUS Variables on page 1172
- For information about Junos OS default groups, see the *CLI User Guide*
- For information about firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

Dynamically Attaching Filters Using RADIUS Variables

You can attach filters to static interfaces by using dynamic profiles. By specifying a variable for the input and output filters, the dynamic profile uses RADIUS VSA attributes for ingress and egress policy.

RADIUS VSA	Attribute Name	Variable
26–10	Ingress-Policy-Name	$junos-input-filter
26–11	Egress-Policy-Name	$junos-output-filter
26–106	IPv6-Ingress-Policy-Name	$junos-input-ipv6-filter
26–107	IPv6-Egress-Policy-Name	$junos-output-ipv6-filter

Before you can attach a filter using RADIUS.

1. Create a basic dynamic profile.

   See "Configuring a Basic Dynamic Profile" on page 673.

2. Ensure that RADIUS ingress and egress policies are configured appropriately.

   See "Configuring RADIUS Server Parameters for Subscriber Access" on page 36.

To dynamically attach IPv4 input and output filters using RADIUS:

1. Specify the dynamic profile you want to attach, the interface, the logical unit number, and family **inet**.

   ```
   [edit]
   user@host# edit dynamic-profiles myProfile interface ge-1/1/1 unit 1 family inet
   ```

2. Specify the IPv4 input filter variable in the dynamic profile.

   ```
   [edit dynamic-profiles myProfile interfaces ge-1/1/1 unit 1 family inet]
   user@host# set filter input $junos-input-filter
   ```

3. Specify the IPv4 output filter variable in the dynamic profile.

   ```
   [edit dynamic-profiles myProfile interfaces ge-1/1/1 unit 1 family inet]
   user@host# set filter output $junos-output-filter
   ```

To dynamically attach IPv6 input and output filters using RADIUS:

1. Specify the dynamic profile you want to attach, the interface, the logical unit number, and family **inet6**.

   ```
   [edit]
   user@host# edit dynamic-profiles myProfile interface ge-1/1/1 unit 1 family inet6
   ```

2. Specify the IPv6 input filter variable in the dynamic profile.

   ```
   [edit dynamic-profiles myProfile interfaces ge-1/1/1 unit 1 family inet6]
   user@host# set filter input $junos-input-ipv6-filter
   ```


Copyright © 2013, Juniper Networks, Inc.

3. Specify the IPv6 output filter variable in the dynamic profile.

   [edit dynamic-profiles myProfile interfaces ge-1/1/1 unit 1 family inet6]
   user@host# set filter output $junos-output-ipv6-filter

Related Documentation

- Classic Filters Overview on page 1133
- Dynamically Attaching Statically Created Filters for Any Interface Type on page 1170
- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type on page 1169
- For more information about Junos default groups, see the *CLI User Guide*
- For more information about firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*

Defining Dynamic Filter Processing Order

You can force filter processing to occur in a particular order by using the **precedence** statement. You specify a precedence for input and output filters within a dynamic profile at the **[edit dynamic-profiles *profile-name* interfaces (*interface-name* | demux0) unit *logical-unit-number* family *family*]** hierarchy level.

The precedence range is from 0 through 250. Setting a lower precedence value for a filter gives it a higher precedence within the dynamic profile. A precedence of zero (the default) gives the filter the highest precedence. If no precedence is specified, the filter receives a precedence of zero (highest precedence). Filters with matching precedence (zero or otherwise) are applied in random order.

Before you define a precedence for a filter in a dynamic profile.

1. Create the filters you want to attach to the dynamic profile.

   See the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices* for detailed information about firewall filters and how to create them.

2. Create a basic dynamic profile.

   See "Configuring a Basic Dynamic Profile" on page 673.

3. Attach the filters to the dynamic profile.

   See "Dynamically Attaching Statically Created Filters for Any Interface Type" on page 1170, "Dynamically Attaching Statically Created Filters for a Specific Interface Family Type" on page 1169, or "Dynamically Attaching Filters Using RADIUS Variables" on page 1172.

To define a precedence for an input and output filter:

1. Specify the input filter precedence in the dynamic profile.

   [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*]
   user@host# set filter input precedence 50

2. Specify the output filter precedence in the dynamic profile.

   [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*]
   user@host# **set filter output precedence 5**

Related Documentation

- Classic Filters Overview on page 1133
- For information about firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*

Configuring Firewall Filter Bypass

You can streamline the filter process, decrease the amount of packet handling for each filter in a chain, and effectively bypass unnecessary filters by using the **service-filter-hit** match/action combination at the **[edit *firewall family family-name* filter *filter-name term term-name*]** hierarchy level.

To bypass firewall filters using the **service-filter-hit** match/action combination, you configure the **service-filter-hit** action in at least one filter in the chain and configure **service-filter-hit** match condition in any subsequent filters that you want to bypass. All packets must pass through each filter in a chain. However, after the **service-filter-hit** flag is set in a packet, the packet "bypasses" any subsequent filters that contain the **service-filter-hit** match condition and more efficiently passes (accepts) marked packets and accelerating the filter process.

NOTE: When using the **service-filter-hit** match/action combination, the order in which the filters are applied is important. You can ensure the order in which the filters are processed by specifying a filter precedence value for the interface. See "Defining Dynamic Filter Processing Order" on page 1173 for more information about dynamic filter processing.

To bypass filter processing:

1. Specify the **service-filter-hit** action for any filters in a filter chain.

   [edit firewall family inet filter video term 1]
   user@host# **set then service-filter-hit**

   When the match conditions for the filter are met, the **service-filter-hit** action is set to indicate to subsequent filters that further processing is unnecessary.

2. Specify the **service-filter-hit** match condition in any filters with a lower precedence (that is, a higher precedence statement value) that you want to detect **service-filter-hit** actions applied from previous filters in the chain.

   [edit firewall family inet filter data term 1]
   user@host# **set from service-filter-hit**

3. Configure the filter to pass (accept) any packet that has a **service-filter-hit** action applied from any previous filters.

Copyright © 2013, Juniper Networks, Inc.

```
[edit firewall family inet filter data term 1]
user@host# set then accept
```

Related Documentation

- Classic Filters Overview on page 1133
- Defining Dynamic Filter Processing Order on page 1173
- Example: Bypassing Firewall Filters on page 1210

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.

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 on page 1133
- Defining Dynamic Filter Processing Order on page 1173
- RADIUS Accounting Statistics for Subscriber Access Overview on page 26
- Configuring Per-Service Session Accounting on page 35
- Configuring Per-Subscriber Session Accounting on page 29
- *Guidelines for Configuring Firewall Filters*
- *Guidelines for Applying Firewall Filters*
- *Firewall Filter Terminating Actions*
- *Firewall Filter Nonterminating Actions*

Configuring a Filter for Use with Enhanced Network Services Mode

For a statically-applied enhanced mode filter to function on statically created interfaces, you must include the **enhanced mode** statement in each filter. However, you do not need to configure the **enhanced mode** statement in filters that are dynamically applied to either static or dynamically-created interfaces.

NOTE: For either static or dynamic interfaces to use enhanced network services mode, you must configure the router chassis network services to use either Enhanced IP Network Services mode or Enhanced Ethernet Network Services mode. By configuring chassis network services to run in one of the enhanced modes, the router enables only MPCs and MS-DPCs in the chassis. See "Firewall Filters and Enhanced Network Services Mode Overview" on page 1141 for details.

To configure a stateless firewall filter to use enhanced mode:

1. Create or edit the stateless firewall filter.

Copyright © 2013, Juniper Networks, Inc.

NOTE: You can configure enhanced mode firewall filters for only inet and inet6 filter families.

For IPv4:

```
[edit]
user@host# edit firewall family inet filter filter-name
```

For IPv6:

```
[edit]
user@host# edit firewall family inet6 filter filter-name
```

2. Specify the filter as an enhanced mode filter.

```
[edit firewall family inet filter filter-name]
user@host# set enhanced-mode
```

3. Configure or modify any filter terms.

   See any of the filter configuration examples described in the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*.

Related Documentation

- *Understanding How to Use Firewall Filters* in the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*
- *Network Services Mode Overview* in the *Junos OS Administration Library for Routing Devices*
- Firewall Filters and Enhanced Network Services Mode Overview on page 1141
- *Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers* in the *Junos OS Administration Library for Routing Devices*
- Dynamic Firewall Filters Overview on page 1132

Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions

Subscriber management enables you to use dynamic profiles to dynamically apply policies that are defined in Ascend-Data-Filters (RADIUS attribute 242) to subscriber sessions. The dynamic profiles include a Junos OS predefined variable that maps the rules and actions defined in the Ascend-Data-Filter to Junos OS features. The RADIUS administrator configures the Ascend-Data-Filter on the RADIUS server in a separate operation.

Subscriber management dynamic profiles use the following Junos OS predefined variables to map family-specific Ascend-Data-Filter rules to Junos OS filter functionality:

- **$junos-adf-rule-v4**—Used for IPv4 family **inet**.
- **$junos-adf-rule-v6**—Used for IPv6 family **inet6**.

To configure a dynamic profile to dynamically apply the policy defined by an Ascend-Data-Filter to a subscriber session:

1. Specify the dynamic profile in which you want to include the Ascend-Data-Filter. Specify the interface, the logical unit number, and the family type.

   [edit]
   user@host# **edit dynamic-profiles** *profile-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *family*

2. Specify that you want to include an Ascend-Data-Filter in the dynamic profile.

   [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*]
   user@host# **edit filter adf**

3. Specify the Junos OS predefined variable that maps the Ascend-Data-Filter actions to Junos OS filter functionality. Use the variable that corresponds to the specified family type.

   [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter adf]
   user@host# **set rule ($junos-adf-rule-v4 | $junos-adf-rule-v6)**

   NOTE: You can also statically configure the Ascend-Data-Filter in this step by entering the filter in hexadecimal format, rather than use a predefined variable. You might use a static filter for testing purposes.

4. (Optional) Suppress error-reporting in the event the RADIUS reply messages do not include the Ascend-Data-Filter attribute.

   [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter adf]
   user@host# **set not-mandatory**

5. (Optional) Enable the counter feature. The counter increments each time a packet matches the rule.

   [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter adf]
   user@host# **set counter**

6. (Optional) Specify the input precedence used to establish the order in which filters on the interface are applied. A lower precedence value equals a higher precedence. The precedence relates to other dynamic filters configured on the same interface.

   [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter adf]
   user@host# **set input-precedence** *precedence*

7. (Optional) Specify the output precedence used to establish the order in which filters on the interface are applied. A lower precedence value equals a higher precedence. The precedence relates to other dynamic filters configured on the same interface.

   [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter adf]
   user@host# **set output-precedence** *precedence*

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Ascend-Data-Filter Policies for Subscriber Management Overview on page 1136
- Ascend-Data-Filter Attribute Fields on page 1138
- Verifying and Managing Dynamic Ascend-Data-Filter Policy Configuration on page 1179
- Example: Configuring Dynamic Ascend-Data-Filter Support for Subscriber Access on page 1202
- Example: Configuring Static Ascend-Data-Filter Support for Subscriber Access on page 1205

Verifying and Managing Dynamic Ascend-Data-Filter Policy Configuration

Purpose View or manage information for Ascend-Data-Filters.

Action

- To display statistics for Ascend-Data-Filters:

  user@host> **show firewall**

- To display firewall log information:

  user@host> **show subscribers extensive**

- To clear filter counters:

  user@host> **clear firewall all**

Related Documentation

- Ascend-Data-Filter Policies for Subscriber Management Overview on page 1136
- Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions on page 1177

Configuring Unicast RPF and Fail Filters in Dynamic Profiles for Subscriber Interfaces

This topic provides a summary of unicast RPF configuration for subscriber interfaces in dynamic profiles on MX Series routers. Unicast RPF provides a way to reduce the effect of denial-of-service attacks on IPv4 and IPv6 interfaces by checking the source IP address against the routing table. Packets that do not match are silently discarded, unless an optional fail filter is configured. The fail filter performs an additional check and directs some action be taken on certain packets. Typical actions include logging the packets or passing them even though they failed the RPF check.

NOTE: **Although the fail filter is technically optional, for dynamic profiles in a DHCP environment you must configure a filter to pass DHCP packets. By default, the RPF check prevents DHCP packets from being accepted on interfaces protected by the RPF check. The fail filter identifies the DHCP packets and passes them on.**

To configure unicast RPF in dynamic profiles:

1. Enable unicast RPF on one or more interfaces in a dynamic profile.

See "Configuring Unicast RPF in Dynamic Profiles for Subscriber Interfaces" on page 1180.

2. (Optional) Create a fail filter to evaluate failed packets and perform further actions.

   See "Configuring a Fail Filter for Unicast RPF in Dynamic Profiles for Subscriber Interfaces" on page 1180.

Related Documentation

- *Unicast RPF in Dynamic Profiles for Subscriber Interfaces*
- Example: Configuring Unicast RPF in a Dynamic Profile on MX Series Routers on page 1219

Configuring Unicast RPF in Dynamic Profiles for Subscriber Interfaces

This topic describes how to configure unicast RPF for subscriber interfaces in dynamic profiles on MX Series routers.

To configure a unicast RPF with a fail filter in a dynamic profile:

1. Access the dynamic profile.

   [edit]
   user@host# **edit dynamic-profiles** ***profile-name***

2. Access the interface and specify the address family

   [edit dynamic-profiles *profile-name*]
   user@host# **edit interfaces** ***interface-name*** **unit** ***logical-unit-number*** **family inet**

3. Configure the RPF check and specify the fail filter.

   [edit dynamic-profiles *profile-name* interface *interface-name* unit *logical-unit-number* family inet]
   user@host# **set rpf-check fail-filter** ***filter-name***

Related Documentation

- Configuring Unicast RPF and Fail Filters in Dynamic Profiles for Subscriber Interfaces on page 1179
- Example: Configuring Unicast RPF in a Dynamic Profile on MX Series Routers on page 1219

Configuring a Fail Filter for Unicast RPF in Dynamic Profiles for Subscriber Interfaces

This topic describes how to configure a fail filter at the **[edit firewall]** hierarchy level that can be optionally applied by unicast RPF for subscriber interfaces in dynamic profiles on MX Series routers.

NOTE: In contrast to statically configured fail filters, RPF-check fail filters used in a dynamic profile cannot be specific to a particular interface.

To configure a firewall fail filter:

1. Create the filter.

Copyright © 2013, Juniper Networks, Inc.

```
[edit]
user@host# edit firewall family inet filter filter-name
```

2. Specify a term for the filter.

```
[edit firewall family inet filter filter-name]
user@host# edit term term-name
```

3. Configure the match conditions for the filter.

```
[edit firewall family inet filter filter-name term term-name]
user@host# set from match-conditions
```

4. Configure the actions to be taken for the matching packets.

```
[edit firewall family inet filter filter-name term term-name]
user@host# set then actions
```

5. (Optional) Repeat Steps 3 and 4 for additional filter terms.

Related Documentation

- Configuring Unicast RPF and Fail Filters in Dynamic Profiles for Subscriber Interfaces on page 1179
- Example: Configuring Unicast RPF in a Dynamic Profile on MX Series Routers on page 1219

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 66

Configuring Fast Update Filters

- Configuring Fast Update Filters on page 1183
- Configuring the Match Order for Fast Update Filters on page 1184
- Configuring Terms for Fast Update Filters on page 1185
- Fast Update Filter Match Conditions on page 1186
- Fast Update Filter Actions and Action Modifiers on page 1187
- Configuring Filters to Permit Expected Traffic on page 1187
- Avoiding Conflicts When Terms Match on page 1188
- Associating Fast Update Filters with Interfaces in a Dynamic Profile on page 1193
- Verifying and Managing Firewall Filter Configuration on page 1194

Configuring Fast Update Filters

You configure a fast update filter in a dynamic profile—this enables you to use dynamic variables in the filter configuration. After you configure fast update filters, you then use the **dynamic-profiles** syntax to associate the filter with the subscriber interface.

To configure a fast update filter for subscriber access:

1. Access the dynamic profile you want to use.

   [edit]
   user@host# **edit dynamic-profiles myProfile**

2. Specify that you want to configure a firewall, and specify the family.

   [edit dynamic-profiles myProfile]
   user@host# **edit firewall family inet**

3. Specify that you want to configure a fast update filter and assign a name to the filter.

   [edit dynamic-profiles myProfile firewall family inet]
   user@host# **edit fast-update-filter httpFilter**

4. Specify the **interface-specific** statement. This statement is mandatory.

   [edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter]
   user@host# **set interface-specific**

5. Configure the match order to use for the filter terms.

Copyright © 2013, Juniper Networks, Inc.

```
[edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter]
user@host# set match-order [source-address protocol destination-port]
```

See "Configuring the Match Order for Fast Update Filters" on page 1184.

6. Specify that you want to configure a term for the filter and assign the name to the term. Configure the match conditions and actions for the term.

```
[edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter]
user@host# edit term term1
```

```
[edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter term
  term1]
user@host# set from protocol tcp
user@host# set from source-address $junos-subscriber-ip-address
user@host# set from destination-port http
user@host# set then count http-cnt
```

See "Configuring Terms for Fast Update Filters" on page 1185.

Related Documentation

- Configuring the Match Order for Fast Update Filters on page 1184
- Configuring Terms for Fast Update Filters on page 1185
- Associating Fast Update Filters with Interfaces in a Dynamic Profile on page 1193
- Fast Update Filters Overview on page 1146
- Dynamic Profiles Overview on page 638
- For information about firewall filters, see *Guidelines for Configuring Firewall Filters* and *Guidelines for Applying Firewall Filters*in the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*.

Configuring the Match Order for Fast Update Filters

You must include the **match-order** statement to explicitly specify the order in which router examines the match conditions. The router examines only those match conditions that you include in the statement. You can match a maximum of five conditions.

NOTE: If the **match-order** statement contains a condition that is not specified in the **from** statement of a term, the router considers that a wildcard for that condition.

If you use the same fast update filter in multiple dynamic profiles, you must configure the same match order for all profiles.

To configure the order in which the router examines the match conditions of a fast update filter:

1. Access the fast update filter:

```
[edit dynamic-profiles myProfile]
```


Copyright © 2013, Juniper Networks, Inc.

```
user@host# edit firewall family inet fast-update-filter httpFilter
```

2. Specify the mandatory **interface-specific** statement.

```
[edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter]
user@host# set interface-specific
```

3. Configure the match order for the match conditions in the filter. Use brackets to enclose multiple match conditions.

```
[edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter]
user@host# set match-order [source-address protocol destination-port]
```

Related Documentation

- Configuring Fast Update Filters on page 1183
- Configuring Terms for Fast Update Filters on page 1185
- Fast Update Filters Overview on page 1146
- Dynamic Profiles Overview on page 638
- Fast Update Filter Match Conditions on page 1186
- For information about firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*

Configuring Terms for Fast Update Filters

A fast update filter consists of one or more terms. A term is made up of one or more match conditions and the action to take when a packet matches the specified conditions.

To configure a term for a fast update filter:

1. Access the fast update filter.

```
[edit dynamic-profiles myProfile]
user@host# edit firewall family inet fast-update-filter httpFilter
```

2. Create the new term and assign a name to the term.

```
[edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter]
user@host# set term term1
```

3. Configure the match condition for the term. See "Fast Update Filter Match Conditions" on page 1186 for the supported match conditions for fast update filters.

```
[edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter]
user@host# set from protocol tcp
user@host# set from source-address $junos-subscriber-ip-address
user@host# set from destination-port http
```

4. Configure the action that the router takes when the match conditions are met. See "Fast Update Filter Actions and Action Modifiers" on page 1187 for the supported actions for fast update filters.

```
[edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter]
user@host# set then accept
```

5. (Optional) Configure the action modifiers that you want the router to take when the match conditions are met. See "Fast Update Filter Actions and Action Modifiers" on page 1187 for the supported action-modifiers for fast update filters.

   [edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter]
   user@host# set then count http-cnt

6. (Optional) Configure the term to be added only once, when the fast update filter is first created.

   [edit dynamic-profiles myProfile firewall family inet fast-update-filter httpFilter]
   user@host# set only-at-create

Related Documentation

- Configuring Fast Update Filters on page 1183
- Configuring the Match Order for Fast Update Filters on page 1184
- Fast Update Filters Overview on page 1146
- Fast Update Filter Match Conditions on page 1186
- Fast Update Filter Actions and Action Modifiers on page 1187
- For additional information about firewall filter terms, see the following topics in the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*
  - *Stateless Firewall Filter Overview*
  - *Stateless Firewall Filter Components*

Fast Update Filter Match Conditions

Table 106: Fast Update Filter Match Conditions

Match Condition	Description
destination-address *prefix*	IP destination address field.
destination-port *number*	TCP or UDP destination port field. Can be a single number, a single range, or one of the standard port synonyms.
dscp *number*	Differentiated services code point. Can be a single number, a single range, or the standard synonyms. IPv4 only.
match-terms *string-of-conditions*	Series of match conditions. Enclose the string within quotation marks and use semicolons to separate entries. For example, **match-terms "protocol tcp; destination-port http";**. Dynamic profile variables are not allowed in the string.
protocol *number*	IP protocol field. Can be a single number, a single range, or one of the standard protocol synonyms. IPv4 only.
source-address *prefix*	IP source address field.
source-port *number*	TCP or UDP source port field. Can be a single number, a single range, or one of the standard protocol synonyms.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring Fast Update Filters on page 1183

Fast Update Filter Actions and Action Modifiers

Table 107: Fast Update Filter Actions and Action Modifiers

Action or Action Modifier	Description
Actions	
accept	Accept the packet.
action-terms *string-of-actions*	A series of multiple actions or action modifiers. Enclose the string within quotation marks and use semicolons to separate entries. For example, **action-terms "log; count http-cnt";**. Dynamic profile variables are not allowed in the string.
discard	Drop the packet silently, without sending an Internet Control Message Protocol (ICMP) message.
ignore-term	Do not add this term to the filter. All match conditions and actions are ignored.
port-mirror	Port mirror packets.
routing-instance *routing-instance*	Forward packets to specified routing instance.
Action Modifiers	
count *counter-name*	Increment the specified counter.
forwarding-class *class*	Classify the packet into one of the following forwarding classes: **as**, **assured-forwarding**, **best-effort**, **expedited-forwarding**, or **network-control**.
log	Log the packet header information.
loss-priority (high | medium-high | medium-low| low)	Set the loss priority level for packets.
policer *policer-name*	Rate-limit packets based on the specified policer.

Related Documentation

- Configuring Fast Update Filters on page 1183

Configuring Filters to Permit Expected Traffic

You must explicitly configure your firewall filter to permit expected traffic, such as DHCP traffic, to pass. Otherwise, the expected traffic is denied when the filter is applied to the interface. This requirement applies to both classic and fast update filters.

The following example shows a fast update filter that might be used to accept DHCP traffic. The actual filter you use depends on the expected traffic in your network.

In the example, the term **allow-dhcp** accepts all DHCP traffic from all source addresses. The term also includes the **only-at-create** option to specify that the term is applied only when the filter is first applied. The term **sub-allow-dhcp** includes the Junos OS predefined variable **$junos-subscriber-ip-address**, which permits all subscriber-specific DHCP traffic.

The **match-order** statement configuration lists the conditions from most-specific to least-specific, as recommended in "Configuring the Match Order for Fast Update Filters" on page 1184. Because this filter is designed to permit ingress DHCP traffic, the **source-address** condition is listed first.

```
firewall {
  family inet {
    fast-update-filter psf1 {
      interface-specific;
      match-order [ source-address destination-address protocol destination-port ];
      term allow-dhcp {
        only-at-create;
        from {
          source-address 0.0.0.0/32;
          destination-address 255.255.255.255/32;
          destination-port 67;
          protocol udp;
        }
        then accept;
      }
      term sub-allow-dhcp {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 192.168.1.2/32;
          destination-port 67;
          protocol udp;
        }
        then accept;
      }
    }
  }
}
```

Related Documentation

- Configuring the Match Order for Fast Update Filters on page 1184
- Configuring Terms for Fast Update Filters on page 1185

Avoiding Conflicts When Terms Match

A fast update filter can contain multiple terms, each with a variety of match conditions. However, when you configure multiple terms in a filter, you must ensure that the terms do not overlap, or conflict with each other. Two terms are considered to overlap when it is possible for a packet to match all conditions of both terms. Because each term specifies a different action for matches, the router cannot determine which action to take. When

 Copyright © 2013, Juniper Networks, Inc.

terms overlap, a conflict error occurs and the session fails when the dynamic profile attempts to apply the filter. The error log indicates the overlapping terms.

How the Router Evaluates Terms in a Filter

The router creates a table of match conditions when examining terms. The table, which is similar to a routing table, is based on the conditions included in the **match-order** statement. When the router receives a packet, the router examines the packet's contents in the sequence specified in the **match-order** statement.

For example, using the sample configuration in the following Match-Order Example, the router first examines the packet's **source-address**, then the **destination-address**, and finally the **destination-port**. As shown in the following table, the two terms in the filter do not overlap because each term has a different **destination-port** specification. The router then takes the appropriate filter action for the term that matches the **destination-port** value of the packet.

Term	source-address	destination-address	destination-port	Action
t55	subscriber's address	3.1.1.2/32	http	count t55_cntr accept
t999	subscriber's address	3.1.1.2/32	https	count t999_cntr accept

Match-Order Example

```
firewall {
  family inet {
    fast-update-filter psf1 {
      interface-specific;
      match-order [ source-address destination-address destination-port ];
      term t55 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 3.1.1.2/32;
          destination-port http;
        }
        then {
          count t55_cntr;
          accept;
        }
      }
      term t999 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 3.1.1.2/32;
          destination-port https;
        }
        then {
          count t999_cntr;
          accept;
        }
      }
```

```
        }
      }
    }
```

Using Implied Wildcards

This section shows an example of how you might use an implied wildcard specification in the match configuration. A condition in the **match-order** statement is an implied wildcard when that condition is not configured in the **from** specification of a term in the filter.

NOTE: When you use ranges (for example, a range of values or a wildcard) in terms, the ranges must not overlap—overlapping ranges create a conflict error. However, you can configure a range in one term and an exact match in another term. For example, in the following filter table, the wildcard destination port value in term **t3** does not overlap the destination port specifications in terms **t55** and **t999** because the **http** and **https** values are exact matches.

In the Implied Wildcard Example configuration, the router views the **destination-port** condition in the **match-order** statement as an implied wildcard for term **t3**, because there is no **destination-port** value configured in that term. As a result, the wildcard specifies that for term **t3** any **destination-port** value is accepted. The filter table appears as follows:

Term	source-address	destination-address	destination-port	Action
t3	subscriber's address	3.1.1.2/32	any (wildcard)	count t3_cntr accept
t55	subscriber's address	3.1.1.2/32	http	count t55_cntr accept
t999	subscriber's address	3.1.1.2/32	https	count t999_cntr accept

In the following filter configuration, traffic with a destination port of **http** matches term **t55** and traffic with a destination port of **https** matches term **t999**. Traffic with a destination port other than **http** or **https** matches term **t3**, which is the implied wildcard.

Implied Wildcard Example

```
firewall {
  family inet {
    fast-update-filter psf1 {
      interface-specific;
      match-order [ source-address destination-address dscp protocol destination-port
        ];
      term t3 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 3.1.1.2/32;
```


Copyright © 2013, Juniper Networks, Inc.

```
            }
            then {
                count t3_cntr;
                accept;
            }
        }
        term t55 {
            from {
                source-address $junos-subscriber-ip-address;
                destination-address 3.1.1.2/32;
                destination-port http;
            }
            then {
                count t55_cntr;
                accept;
            }
        }
        term t999 {
            from {
                source-address $junos-subscriber-ip-address;
                destination-address 3.1.1.2/32;
                destination-port https;
            }
            then {
                count t999_cntr;
                accept;
            }
        }
    }
  }
}
```

Conflict Caused by Overlapping Ranges

This section shows two examples of overlapping ranges in terms. When you use ranges (such as a wildcard or a range of values) in terms, the ranges must not overlap—overlapping ranges create a conflict error and the session fails.

In the following filter configuration, the **destination-port** ranges in the two terms overlap. Ports in the range from 50 through 80 match both term **src0** and term **src1**, which each specify different actions to take.

NOTE: You can configure a range in one term and an exact match in another term. See the section, *Using Implied Wildcards*, for an example that uses a wildcard for a match condition in one term and an exact match for the condition in a second term.

Term	source-address	destination-address	destination-port	Action
src0	subscriber's address	10.1.1.2/32	0–80	count c1_cntr accept

Term	source-address	destination-address	destination-port	Action
src1	subscriber's address	10.1.1.2/32	50–100	count c2_cntr accept

Overlapping Ranges Example 1

```
firewall {
  family inet {
    fast-update-filter fuf–src {
      interface-specific;
      match-order [ source-address destination-address destination-port ];
      term src0 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 10.1.1.2/32;
          destination-port 0–80;
        }
        then {
          count c1_cntr;
          accept;
        }
      }
      term src1 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 10.1.1.2/32;
          destination-port 50–100;
        }
        then {
          count c2_cntr;
          accept;
        }
      }
    }
```

In this filter configuration, the **protocol** specification in terms **src21** and **src22** use the implied wildcard, which configures a range for each term. Because overlapping ranges are not allowed, a conflict error results.

Term	source-address	destination-address	protocol	destination-port	Action
src20	subscriber's address	10.1.1.2/32	udp	any (wildcard)	count c20_cntr accept
src21	subscriber's address	10.1.1.2/32	any (wildcard)	http	count c21_cntr accept
src21	subscriber's address	10.1.1.2/32	any (wildcard)	https	count c22_cntr accept

Copyright © 2013, Juniper Networks, Inc.

Overlapping Ranges Example 2

```
firewall {
  family inet {
    fast-update-filter fuf–src2 {
      interface-specific;
      match-order [ source-address destination-address protocol destination-port ];
      term src20 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 10.1.1.2/32;
          protocol udp;
        }
        then {
          count c20_cntr;
          accept;
        }
      }
      term src21 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 10.1.1.2/32;
          destination-port http;
        }
        then {
          count c21_cntr;
          accept;
        }
      }
      term src22 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 10.1.1.2/32;
          destination-port https;
        }
        then {
          count c22_cntr;
          accept;
        }
      }
    }
```

Related Documentation

- Configuring Fast Update Filters on page 1183
- Configuring Terms for Fast Update Filters on page 1185
- Configuring the Match Order for Fast Update Filters on page 1184

Associating Fast Update Filters with Interfaces in a Dynamic Profile

After you configure the fast update filter, you reference the filter in the **interfaces** stanza of a dynamic profile. When the dynamic profile instantiates a subscriber session, the router applies the terms of the filter to the interface.

To apply a fast update filter to an interface in a dynamic profile:

1. Access the dynamic profile you want to use.

   [edit]
   user@host# **edit dynamic-profiles myProfile**

2. Specify the interface for the dynamic profile—use the dynamic interface variable.

   [edit dynamic-profiles myProfile]
   user@host# **edit interfaces $junos-interface-ifd-name**

3. Specify the underlying interface—use the unit number variable.

   [edit dynamic-profiles myProfile interfaces “$junos-interface-ifd-name”]
   user@host# **edit unit $junos-underlying-interface-unit**

4. Specify the family. Use **inet** if you are using IPv4 filters or **inet6** for IPv6 filters.

   [edit dynamic-profiles myProfile interfaces “$junos-interface-ifd-name” unit
   “$junos-underlying-interface-unit”]
   user@host# **edit family inet**

5. Specify the filters that you want to apply to the interface.

   [edit dynamic-profiles myProfile interfaces “$junos-interface-ifd-name” unit
   “$junos-underlying-interface-unit” family inet]
   user@host# **set filter input httpFilter**
   user@host# **set filter output myOutFilter**

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring Static Subscriber Interfaces in Dynamic Profiles on page 765
- Associating Dynamic Profiles with Statically Created Interfaces on page 768
- Fast Update Filters Overview on page 1146
- For information about firewall filters, see *Guidelines for Configuring Firewall Filters* and *Guidelines for Applying Firewall Filters*in the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*.

Verifying and Managing Firewall Filter Configuration

Purpose View or manage information for firewall filters:

NOTE: The router creates unique names for fast update filters and for filter terms and counters. See *Naming Fast Update Filters* in “Fast Update Filters Overview” on page 1146 for information.

Action

- To display statistics for firewall filters:

  user@host> **show firewall**

- To display firewall log information:

 Copyright © 2013, Juniper Networks, Inc.

```
user@host> show firewall log
```

- To clear filter counters:

```
user@host> clear firewall all
```

Related Documentation

- Classic Filters Overview on page 1133
- Fast Update Filters Overview on page 1146
- *Junos OS Operational Mode Commands*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 67

Configuring Service Sets in Dynamic Profiles

- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197
- Verifying and Managing Service Sets Information on page 1198

Associating Service Sets with Interfaces in a Dynamic Profile

After you configure a service set, you use a dynamic profile to dynamically associate the service set with interfaces. You reference the filter in the **interfaces** stanza of a dynamic profile. When the dynamic profile instantiates a subscriber session, the router applies the terms of the filter to the interface.

To apply a service set to an interface in a dynamic profile:

1. Access the dynamic profile you want to use.

   ```
   [edit]
   user@host# edit dynamic-profiles myProfile
   ```

2. Specify the interface for the dynamic profile—use the dynamic interface variable.

   ```
   [edit dynamic-profiles myProfile]
   user@host# edit interfaces $junos-interface-ifd-name
   ```

3. Specify the underlying interface—use the unit number variable.

   ```
   [edit dynamic-profiles myProfile interfaces "$junos-interface-ifd-name"]
   user@host# edit unit $junos-underlying-interface-unit
   ```

4. Specify the family. Dynamic service sets are supported only on **family inet** (IPv4).

   ```
   [edit dynamic-profiles myProfile interfaces "$junos-interface-ifd-name" unit
     "$junos-underlying-interface-unit"]
   user@host# edit family inet
   ```

5. Specify the input and output service sets that you want to apply to the interface.

   ```
   [edit dynamic-profiles myProfile interfaces "$junos-interface-ifd-name" unit
     "$junos-underlying-interface-unit" family inet]
   user@host# set service input service-set inputService_200
   user@host# set service input post-service-filter postService_15
   user@host# set service output service-set outputService_320
   ```

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Dynamic Service Sets Overview on page 1151
- Verifying and Managing Service Sets Information on page 1198
- For information about creating service sets, see "Service Set Configuration Guidelines" in the *Junos OS Services Interfaces Library for Routing Devices*.
- For information about statically applying service sets to interfaces, see *Applying Filters and Services to Interfaces* in the *Junos OS Services Interfaces Library for Routing Devices*.

Verifying and Managing Service Sets Information

Purpose View information for service sets:

Action

- To display summary information for service sets:

  user@host> **show services service-sets summary**

- To display interface-specific information for service sets:

  user@host>**show services service-sets summary interface** ***interface-name***

Related Documentation

- Dynamic Service Sets Overview on page 1151
- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197
- *Junos OS Operational Mode Commands*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 68

Firewall Filter Examples

- Examples: Configuring Static Filters on page 1199
- Example: Configuring Dynamic Ascend-Data-Filter Support for Subscriber Access on page 1202
- Example: Configuring Static Ascend-Data-Filter Support for Subscriber Access on page 1205
- Example: Configuring Fast Update Filters for Subscriber Access on page 1209
- Example: Bypassing Firewall Filters on page 1210
- Example: Configuring Hierarchical Policers as Filter Actions on page 1214
- Example: Interface-Shared Filter Configuration on page 1217
- Example: Configuring Unicast RPF in a Dynamic Profile on MX Series Routers on page 1219

Examples: Configuring Static Filters

This topic provides some static filter configuration examples.

```
firewall {
  policer p1 {
    if-exceeding {
      bandwidth-limit 5m;
      burst-size-limit 10m;
    }
    then discard;
  }
  family inet {
    filter dfwd {
      interface-specific;
      term 1 {
        from {
          source-address {
            192.1.1.0/24;
          }
        }
        then {
          count c1;
          next term;
        }
      }
      term 2 {
```

```
        from {
            source-address {
                192.2.1.0/24;
            }
        }
        then count c2;
    }
    term 3 {
        then accept;
    }
}
filter dfwd1 {
    interface-specific;
    term 1 {
        from {
            address {
                192.1.1.0/24;
            }
        }
        then {
            discard;
        }
    }
}
filter tos {
    interface-specific;
    term 1 {
        from {
            precedence priority;
        }
        then forwarding-class assured-forwarding;
    }
    term 2 {
        then {
            log;
            accept;
        }
    }
}
filter dfwd2 {
    interface-specific;
    term 1 {
        from {
            forwarding-class best-effort;
        }
        then {
            sample;
            forwarding-class expedited-forwarding;
        }
    }
    term 2 {
        then accept;
    }
}
filter nodhcp {
    term dhcpdiscover {
```

 Copyright © 2013, Juniper Networks, Inc.

```
            from {
                protocol udp;
                source-port 68;
                destination-port 67;
            }
            then {
                discard;
            }
        }
        term others {
            then accept;
        }
    }
    filter p1 {
        interface-specific;
        term 1 {
            from {
                precedence priority;
            }
            then {
                policer p1;
                log;
            }
        }
        term 2 {
            then accept;
        }
    }
    filter dscp {
        interface-specific;
        term 1 {
            from {
                dscp af11;
            }
            then log;
        }
        term 2 {
            then accept;
        }
    }
    filter tcm {
        interface-specific;
        term 1 {
            from {
                dscp af11;
            }
            then policer p1;
        }
        term 2 {
            then accept;
        }
    }
}
traceoptions {
    flag dynamic;
}
```

```
}
```

Related Documentation

- Dynamically Attaching Statically Created Filters for Any Interface Type on page 1170
- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type on page 1169

Example: Configuring Dynamic Ascend-Data-Filter Support for Subscriber Access

This example shows how to configure support for dynamic Ascend-Data-Filter policies.

- Requirements on page 1202
- Overview on page 1202
- Configuration on page 1202
- Verification on page 1203

Requirements

- Ensure that the Ascend-Data-Filter has been configured on the RADIUS server.
- Create the dynamic profile. See "Dynamic Profiles Overview" on page 638.
- Configure RADIUS support. See "Configuring RADIUS Server Parameters for Subscriber Access" on page 36.

Overview

Ascend-Data-Filters are configured on a RADIUS server, and contain rules that create policies. Subscriber management uses a dynamic profile to obtain the Ascend-Data-Filter attribute (RADIUS attribute 242) from the RADIUS server and apply the policy to a subscriber session.

- Specify the dynamic profile to use to apply the Ascend-Data-Filter policy to the subscriber session.
- Specify the Junos OS predefined variable that maps the Ascend-Data-Filter rules to Junos OS filter functionality.
- Configure optional settings, which include counting the rule usage and setting the precedence order for the filter.

Configuration

Step-by-Step Procedure

To configure dynamic Ascend-Data-Filter support:

1. Specify the dynamic profile in which you want to include the Ascend-Data-Filter, and configure the interface, the logical unit number, and the family type.

```
[edit]
user@host# edit dynamic-profiles adf-profile-v4 interfaces
  $junos-interface-ifd-name unit $junos-underlying-interface-unit family inet
```


Copyright © 2013, Juniper Networks, Inc.

2. Specify that you want to include an Ascend-Data-Filter in the dynamic profile and provide the Junos OS predefined variable as the rule that maps the Ascend-Data-Filter actions to Junos OS filter functionality.

```
[edit dynamic-profiles adf-profile-v4 interfaces “$junos-interface-ifd-name” unit
   “$junos-underlying-interface-unit” family inet]
user@host# set filter adf rule $junos-adf-rule-v4
```

3. Enable the counter for the rule.

```
[edit dynamic-profiles adf-profile-v4 interfaces “$junos-interface-ifd-name” unit
   “$junos-underlying-interface-unit” family inet]
user@host# set filter adf counter
```

4. Specify the precedence for received packets on the interface.

```
[edit dynamic-profiles adf-profile-v4 interfaces “$junos-interface-ifd-name” unit
   “$junos-underlying-interface-unit” family inet]
user@host# set filter adf input-precedence 75
```

5. Specify the precedence for transmitted packets on the interface.

```
[edit dynamic-profiles adf-profile-v4 interfaces “$junos-interface-ifd-name” unit
   “$junos-underlying-interface-unit” family inet]
user@host# set filter adf output precedence 80
```

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

```
[edit]
user@host# show dynamic-profiles
...
adf-profile-v4 {
   interfaces {
      "$junos-interface-ifd-name" {
         unit "$junos-underlying-interface-unit" {
            family inet {
               filter {
                  adf {
                     rule "$junos-adf-rule-v4";
                     counter;
                     input-precedence 75;
                     output-precedence 80;
                  ...
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying that Dynamic Ascend-Data-Filter Rules Are Applied to Subscriber Sessions on page 1204
- Verifying Dynamic Ascend-Data-Filter Usage on page 1205

Verifying that Dynamic Ascend-Data-Filter Rules Are Applied to Subscriber Sessions

Purpose Verify that the Ascend-Data-Filter rules were attached to the subscriber.

Action From operational mode, enter the **show subscribers extensive** command.

```
user@host>show subscribers extensive
  Type: DHCP
  User Name: user1-adf
  IP Address: 192.168.1.10
  IP Netmask: 255.255.255.0
  Logical System: default
  Routing Instance: default
  Interface: ge-1/0/0.0
  Interface type: Static
  Dynamic Profile Name: adf-profile-v4
  MAC Address: 00:10:94:00:00:01
  State: Active
  Radius Accounting ID: 5
  Login Time: 2010-08-12 14:06:27 PDT
  ADF IPv4 Input Filter Name: __junos_adf_5-ge-1/0/0.0-inet-in
          Rule 0: 0101010000000000d87f920000180000000000000000000000
              from {
                destination-address 216.127.146.0/24;
              }
              then {
                accept;
              }
          Rule 1: 0100010000000000000000000000000600000000001900020000
              from {
                protocol 6;
                destination-port 25;
              }
              then {
                discard;
              }
          Rule 2: 0101010000000000000000000000000000000000000000000000
              then {
                accept;
              }
```

Meaning The output shows the information for the dynamic profile, including Ascend-Data-Filter rules. Verify the following information:

- The User Name field indicates the correct subscriber.
- The Dynamic Profile Name field is correct for the subscriber.
- The correct Ascend-Data-Filter rules are applied to the subscriber. The display shows the rules that are configured on the RADIUS server.

 Copyright © 2013, Juniper Networks, Inc.

Verifying Dynamic Ascend-Data-Filter Usage

Purpose Verify usage of the dynamic Ascend-Data-Filter. Counter statistics are displayed when the **counter** option is configured for the **adf** command in the dynamic profile.

Action From operational mode, enter the **show firewall** command.

```
user@host> show firewall

Filter: __junos_adf_5-ge-1/0/0.0-inet-in
Counters:
Name                          Bytes              Packets
t0-cnt                        32758                   22
t1-cnt                        22199                   15
t2-cnt                        21723                   14
```

Meaning The output shows the name of the filter and lists the counter activity. If the **counter** option is not configured, the output displays only the filter name.

Related Documentation

- Ascend-Data-Filter Policies for Subscriber Management Overview on page 1136
- Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions on page 1177

Example: Configuring Static Ascend-Data-Filter Support for Subscriber Access

This example shows how to configure support for static Ascend-Data-Filter policies. In a static configuration, you manually configure the Ascend-Data-Filter as part of the dynamic profile configuration. This procedure differs from dynamic configuration, in which the Ascend-Data-Filter is defined on the RADIUS server and then subscriber management uses a predefined variable to map the Ascend-Data-Filter rules to Junos OS filter functionality. Because creating a static Ascend-Data-Filter configuration can be labor-intensive, you might typically use this method for testing purposes.

- Requirements on page 1205
- Overview on page 1206
- Configuration on page 1206
- Verification on page 1208

Requirements

- Create the dynamic profile. See "Dynamic Profiles Overview" on page 638.
- Configure RADIUS support. See "Configuring RADIUS Server Parameters for Subscriber Access" on page 36.

Overview

Ascend-Data-Filters contain rules that create policies. Subscriber management uses a dynamic profile to apply the policy to a subscriber session. You manually configure the Ascend-Data-Filter as part of the dynamic policy.

- Specify the dynamic profile to use to apply the Ascend-Data-Filter policy to the subscriber session.
- Configure the Ascend-Data-Filter.
- Configure optional settings, which include counting the rule usage and setting the precedence for received and transmitted traffic.

Configuration

Step-by-Step Procedure

To configure static Ascend-Data-Filter support:

1. Specify the dynamic profile in which you want to create the Ascend-Data-Filter, and configure the interface, the logical unit number, and the family type.

```
[edit]
user@host# edit dynamic-profiles adf-profile-v4 interfaces
  $junos-interface-ifd-name unit $junos-underlying-interface-unit family inet
```

2. Configure the Ascend-Data-Filter. Enclose the filter values within quotation marks. You can configure multiple Ascend-Data-Filter rules in the same dynamic profile.

```
[edit dynamic-profiles adf-profile-v4 interfaces "$junos-interface-ifd-name" unit
  "$junos-underlying-interface-unit" family inet]
user@host# set filter adf rule "01000100 0A020100 00000000 18000000
  00000000 00000000"
```

3. Enable the counter for the rule.

```
[edit dynamic-profiles adf-profile-v4 interfaces "$junos-interface-ifd-name" unit
  "$junos-underlying-interface-unit" family inet]
user@host# set filter adf counter
```

4. Specify the precedence for received packets on the interface.

```
[edit dynamic-profiles adf-profile-v4 interfaces "$junos-interface-ifd-name" unit
  "$junos-underlying-interface-unit" family inet]
user@host# set filter adf input-precedence 80
```

5. Specify the precedence for transmitted packets on the interface.

```
[edit dynamic-profiles adf-profile-v4 interfaces "$junos-interface-ifd-name" unit
  "$junos-underlying-interface-unit" family inet]
user@host# set filter adf output precedence 85
```

Results

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

```
[edit]
user@host# show dynamic-profiles
```


Copyright © 2013, Juniper Networks, Inc.

```
...
adf-profile-v4 {
   interfaces {
      "$junos-interface-ifd-name" {
         unit "$junos-underlying-interface-unit" {
            family inet {
               filter {
                  adf {
                     rule "01000100 0A020100 00000000 18000000 00000000 00000000";
                     counter;
                     input-precedence 80;
                     output-precedence 85;
                  ...
```

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

Results

The Ascend-Data-Filter rule defined in Step 2 of the procedure configures an input policy that filters all packets from network 10.2.1.0 with wildcard mask 255.255.255.0 to any destination.

Table 108 on page 1207 lists the values specified in the Ascend-Data-Filter rule.

Table 108: Ascend-Data-Filter Rule

Action or Classifier	Hex Value	Junos OS Filter Function
Type	01	IPv4
Forward	00	Forward
Indirection	01	Ingress
Spare	00	None
Source IP address	0a020100	10.2.1.0
Destination IP address	00000000	Any
Source IP mask	18	24 (255.255.255.0)
Destination IP mask	00	0 (0.0.0.0)
Protocol	00	None
Established	00	None
Source port	0000	None
Destination port	0000	None
Source port qualifier	00	None

Copyright © 2013, Juniper Networks, Inc.

Table 108: Ascend-Data-Filter Rule *(continued)*

Action or Classifier	Hex Value	Junos OS Filter Function
Destination port qualifier	00	None
Reserved	0000	None

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying that Static Ascend-Data-Filter Rules are Applied to Subscriber Sessions on page 1208
- Verifying Static Ascend-Data-Filter Usage on page 1209

Verifying that Static Ascend-Data-Filter Rules are Applied to Subscriber Sessions

Purpose Verify that the Ascend-Data-Filter rules you manually configured were attached to the subscriber.

Action From operational mode, enter the **show subscribers extensive** command.

```
user@host>show subscriber extensive
  Type: DHCP
  User Name: user1-adf
  IP Address: 192.168.1.10
  IP Netmask: 255.255.255.0
  Logical System: default
  Routing Instance: default
  Interface: ge-1/0/0.0
  Interface type: Static
  Dynamic Profile Name: adf-profile-v4
  MAC Address: 00:10:94:00:00:01
  State: Active
  Radius Accounting ID: 5
  Login Time: 2010-08-12 14:06:27 PDT
  ADF IPv4 Input Filter Name: __junos_adf_5-ge-1/0/0.0-inet-in
            Rule 0: 010001000A020100000000001800000000000000000000000
                from {
                  destination-address 10.2.1.0/24;
                }
                then {
                  accept;
                }
```

Meaning The output shows the information for the dynamic profile, including Ascend-Data-Filter rules. Verify the following information:

- The User Name field indicates the correct subscriber.
- The Dynamic Profile Name field is correct for the subscriber.

 Copyright © 2013, Juniper Networks, Inc.

- The correct static Ascend-Data-Filter rule is applied to the subscriber.

Verifying Static Ascend-Data-Filter Usage

Purpose Verify usage of the static Ascend-Data-Filter. Counter statistics are displayed when the **counter** option is configured for the **adf** command in the dynamic profile.

Action From operational mode, enter the **show firewall** command.

```
user@host> show firewall

Filter: __junos_adf_5-ge-1/0/0.0-inet-in
Counters:
Name                          Bytes              Packets
t0-cnt                        32758                   22
```

Meaning The output shows the name of the filter and the lists counter activity. If the **counter** option is not configured, the output displays only the filter name.

Related Documentation
- Ascend-Data-Filter Policies for Subscriber Management Overview on page 1136
- Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions on page 1177

Example: Configuring Fast Update Filters for Subscriber Access

This example shows you how to configure a fast update filter that is an input filter that counts the HTTP and non-HTTP packets from a subscriber. In the example, you use the firewall stanza to create the filter and the interfaces stanza to attach the filter.

```
[edit dynamic-profiles myProfile]
firewall {
  family inet {
    fast-update-filter httpFilter {
      interface-specific;
      match-order [source-address protocol destination-port];
      term term1 {
        from {
          protocol tcp;
          source-address $junos-subscriber-ip-address;
          destination-port http;
        }
        then {
          count http-cnt;
        }
      }
      term term2 {
        from {
          protocol tcp;
          source-address $junos-subscriber-ip-address;
        }
        then {
          count non-http-cnt;
        }
```

```
            }
          }
        }
      }
      interfaces {
        "$junos-interface-ifd-name" {
          unit "$junos-underlying-interface-unit" {
            family inet {
              filter {
                input httpFilter;
              }
            }
          }
        }
      }
```

Related Documentation

- Configuring Fast Update Filters on page 1183

Example: Bypassing Firewall Filters

This example describes how to configure multiple filters using the **service-filter-hit** match/action combination and contains the following sections:

- Before You Begin on page 1210
- Filter Bypass Overview on page 1210
- Configuring Filter Bypass on page 1211

Before You Begin

When using the **service-filter-hit** match/action combination, keep the following in mind:

- The order in which the filters are applied is important. You can ensure the order in which the filters are processed by specifying a filter precedence value for the interface. See "Defining Dynamic Filter Processing Order" on page 1173 for more information about dynamic filter processing and how to use the **precedence** statement.
- The following example uses policers to further define the match conditions each filter uses. These filters are not described here. To better understand how to configure policers, see "*Statement Hierarchy for Configuring Policers*" in the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*.

Filter Bypass Overview

Packets must pass through each filter in a chain. However, if you create a chain of filters to process different types of packets (for example, voice, video, and data packets), you can streamline the filter process, decreasing the amount of packet handling for each filter in the chain, effectively bypassing unnecessary filters, by using the **service-filter-hit** match/action combination at the **[edit** ***firewall family family-name*** **filter** ***filter-name term term-name*****]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

Figure 37 on page 1211 shows the logical processing flow through a chain of three filters (voice, video, and data) where only processing for a specific data type is desired. This configuration example shows an ingress filter flow. Though subsequent ingress filters in a chain can detect whether the **service-filter-hit** action is set, egress filters do not. To bypass egress filters, you must also configure the **service-filter-hit** match/action combination on those filters.

Figure 37: Logical Flow Example for Filter Bypass Processing

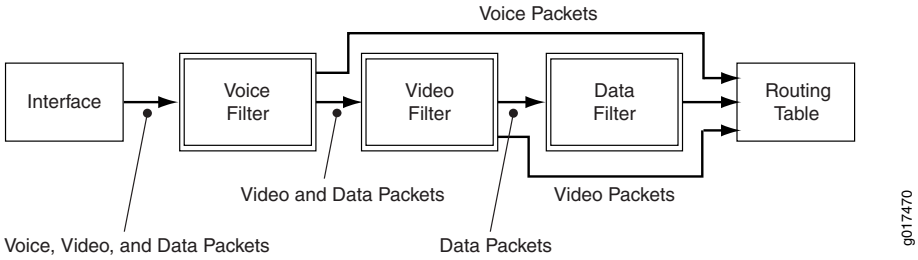

Configuring Filter Bypass

- Configuring the Voice Filter on page 1211
- Configuring the Video Filter on page 1212
- Configuring the Data Filter on page 1212
- Results on page 1212

CLI Quick Configuration

To quickly configure this example:

```
[edit]
set firewall filter voice term T1 from address 1.1.1.1/32
set firewall filter voice term T1 from source-port 5004-5005
set firewall filter voice term T1 then forwarding-class assured-forwarding service-filter-hit
  accept
set firewall filter voice term default then accept
set firewall filter video term T1 from service-filter-hit
set firewall filter video term T1 then accept
set firewall filter video term T2 from source-address 10.10.10.10/32
set firewall filter video term T2 then policer video-policer service-filter-hit accept
set firewall filter video term default then accept
set firewall filter data term T1 from service-filter-hit
set firewall filter data term T1 then accept
set firewall filter data term T2 then policer data-policer service-filter-hit accept
```

Configuring the Voice Filter

Step-by-Step Procedure

To configure the voice filter for the logical flow in Figure 37 on page 1211:

1. Configure the filter to apply the assured forwarding class and set the **service-filter-hit** action for traffic from a specific address and port range (over which voice traffic is expected).

```
[edit]
set firewall filter voice term T1 from address 1.1.1.1/32
set firewall filter voice term T1 from source-port 5004-5005
```

```
set firewall filter voice term T1 then forwarding-class assured-forwarding
  service-filter-hit accept
```

2. Configure the filter default action to pass (accept) packet traffic from any other address or port range.

```
[edit]
set firewall filter voice term default then accept
```

Configuring the Video Filter

Step-by-Step Procedure

To configure the video filter for the logical flow in Figure 37 on page 1211:

1. Configure the filter to pass (accept) incoming packets that are tagged by the **service-filter-hit** action.

```
[edit]
set firewall filter video term T1 from service-filter-hit
set firewall filter video term T1 then accept
```

2. Configure the filter to apply a video policer and set the **service-filter-hit** action for traffic from a specific address (over which video traffic is expected).

```
[edit]
set firewall filter video term T2 from source-address 10.10.10.10/32
set firewall filter video term T2 then policer video-policer service-filter-hit accept
```

3. Configure the filter default action to pass (accept) packet traffic from any other address or port range.

```
[edit]
set firewall filter video term default then accept
```

Configuring the Data Filter

Step-by-Step Procedure

To configure the data filter for the logical flow in Figure 37 on page 1211:

1. Configure the filter to pass (accept) incoming packets that are tagged by the **service-filter-hit** action.

```
[edit]
set firewall filter data term T1 from service-filter-hit
set firewall filter data term T1 then accept
```

2. Configure the filter to apply a data policer and set the **service-filter-hit** action for traffic from a specific address (over which video traffic is expected).

```
[edit]
set firewall filter data term T2 then policer data-policer service-filter-hit accept
```

Results

Display the results of the configuration:

```
[edit firewall]
user@host# show
filter voice {
```


Copyright © 2013, Juniper Networks, Inc.

```
    term T1 {
      from {
        address {
          1.1.1.1/32;
        }
        source-port 5004-5005;
      }
      then {
        forwarding-class assured-forwarding;
        service-filter-hit;
        accept;
      }
    }
    term default {
      then accept;
    }
  }
  filter video {
    term T1 {
      from {
        service-filter-hit;
      }
      then accept;
    }
    term T2 {
      from {
        source-address {
          10.10.10.10/32;
        }
      }
      then {
        policer video_policer;
        service-filter-hit;
        accept;
      }
    }
    term default {
      then accept;
    }
  }
  filter data {
    term T1 {
      from {
        service-filter-hit;
      }
      then accept;
    }
    term T2 {
      then {
        policer data_policer;
        service-filter-hit;
        accept;
      }
    }
  }
```

Related Documentation

- Classic Filters Overview on page 1133
- Defining Dynamic Filter Processing Order on page 1173
- *Statement Hierarchy for Configuring Policers*
- Configuring Firewall Filter Bypass on page 1174

Example: Configuring Hierarchical Policers as Filter Actions

This example shows how to configure a hierarchical policer and apply the policer to ingress Layer 3 traffic at a logical interface on the MX-series platform.

- Requirements on page 1214
- Overview on page 1214
- Configuration on page 1214
- Verification on page 1217

Requirements

Before you begin, be sure that your environment meets the following requirements:

- Supported on MX Series routers.

Overview

In this example, you configure a hierarchical policer as a filter action.

Configuration

- Example: Hierarchical Policer as Filter Action on page 1214
- Example: Defining the Interface: on page 1216

Example: Hierarchical Policer as Filter Action

Step-by-Step Procedure

You can have hierarchical policers as one type of filter action. To configure a firewall filter:

1. Configure the family address type for a firewall filter:

   ```
   [edit firewall]
   user@host# set family inet
   ```

2. Specify the filter name:

   ```
   [edit firewall family inet]
   user@host# set filter inet-filter
   ```

3. Specify the term name:

   ```
   [edit firewall family inet filter inet-filter]
   user@host# set term t1
   ```


Copyright © 2013, Juniper Networks, Inc.

4. In each firewall filter term, specify the match conditions to use to match components of a packet:

```
[edit firewall family inet filter inet-filter term t1]
user@host# set from precedence critical-ecp immediate priority
user@host# set from protocol tcp
```

5. In each firewall filter term, specify the actions to take if the packet matches all the condition in that term:

```
[edit firewall family inet filter inet-filter term t1]
user@host# set then hierarchical-policer HP1
```

6. (Optional) Enable all hierarchical policers in one filter to share the same policer instance in PFE:

```
[edit firewall family inet filter inet-filter term t1]
user@host# set then hierarchical-policer HP1 filter-specific
```

Results Confirm the configuration by entering the **show firewall** configuration command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
   filter inet-filter {
      interface-specific;
      term t1 {
         from {
            precedence [ critical-ecp immediate priority ];
            protocol tcp;
         }
         then hierarchical-policer HP1;
      }
      term t2 {
         from {
            precedence [ internet-control routine ];
            protocol tcp;
         }
         then hierarchical-policer HP2;
      }
   }
}
family inet6 {
   filter inet6-filter {
      interface-specific;
      term t1 {
         from {
            next-header [ tcp udp ];
            forwarding-class [ assured-forwarding expedited-forwarding ];
         }
         then hierarchical-policer HP1;
      }
      term t2 {
         from {
            next-header [ tcp udp icmpv6 ospf rsvp ];
```

```
            forwarding-class [ network-control best-effort ];
          }
          then hierarchical-policer HP2;
        }
      }
    }
  }
```

Example: Defining the Interface:

Step-by-Step Procedure

To define the interface:

1. Enable configuration of the physical interface:

   ```
   [edit]
   user@host# edit interfaces ge-1/2/0 unit 0
   ```

2. Configure the family address:

   ```
   [edit interfaces ge-1/2/0 unit 0]
   user@host# set family inet address 10.100.16.2/24
   ```

3. Specify the filter name:

   ```
   [edit interfaces ge-1/2/0 unit 0 family inet]
   user@host# set filter inet-filter
   user@host# set address 10.100.16.2/24
   ```

Results

Confirm the configuration by entering the **show interfaces** configuration command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
interfaces {
  ge-1/2/0 {
    unit 0 {
      family inet {
        filter {
          input inet-filter;
        }
        address 10.100.16.2/24;
      }
      family inet6 {
        input-hierarchal-policer shared_HP;
        address 1A23:120B::7634:AD01:4D/120;
      }
    }
  }
  ge-1/2/1 {
    unit 0 {
      family inet {
        input-hierarchal-policer shared_HP;
        address 10.100.16.2/24;
      }
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
```

Verification

Confirm that the configuration is working properly.

- Displaying Packets for the Firewall on page 1217

Displaying Packets for the Firewall

Purpose Verify the number of packets evaluated by the policer. Premium policer counters are not supported.

Action Use the **show firewall** operational mode command. The command output displays the number of packets.

```
[edit]
user@host# show firewall
```

```
Filter: __default_bpdu_filter__

Filter: utp_4550-ge-1/0/0.100-in
Counters:
Name                                                  Bytes              Packets
c_ef-ge-1/0/0.0-i                                   1696750                15425
c_other-ge-1/0/0.0-i                                      0                    0
Policers:
Name                                                Packets
hp_abc-filter-ge-1/0/0.0-i                             7509
```

Related Documentation

- Hierarchical Policer Overview on page 1142
- Hierarchical Policer as Filter Action on page 1143
- filter-specific on page 1642

Example: Interface-Shared Filter Configuration

Before you can attach an interface-shared filter using a dynamic profile.

- Create a basic dynamic profile.

  See "Configuring a Basic Dynamic Profile" on page 673.

To configure an interface-shared filter using a dynamic profile that is used to implement agent-circuit-identifier VLAN household filtering:

1. Access the dynamic profile you want to use.

   [edit]
   user@host# edit dynamic-profiles *profile-name*

2. Specify the interfaces.

   [edit dynamic-profiles *profile-name*]

user@host# **edit** interfaces *interface-name*

3. Specify the unit.

   [edit dynamic-profiles *profile-name* interfaces *interface-name*]
   user@host# **edit unit $junos-interface-unit**

4. Specify the family.

   [edit dynamic-profiles *profile-name* interfaces *interface-name* unit
   "$junos-interface-unit"]
   user@host# **edit family** *family-name*

5. Specify the input filter and the filter terms for the interface unit.

   [edit dynamic-profiles *profile-name* interfaces *interface-name* unit
   "$junos-interface-unit" family *family-name*]
   user@host# **edit input $junos-input-filter shared-name $junos-interface-set-name**
   precedence *precedence-number*

6. Specify the output filter and the filter terms for the interface unit.

   [edit dynamic-profiles *profile-name* interfaces*interface-name* unit
   "$junos-interface-unit" family *family-name*]
   user@host# **edit input $junos-output-filter shared-name $junos-interface-set-name**
   precedence *precedence-number*

7. Specify that you want to configure a firewall, and specify the family.

   [edit dynamic-profiles *profile-name*]
   user@host# **edit firewall family** *family-name*

8. Specify the filter.

   [edit dynamic-profiles *profile-name* firewall family *family-name*
   user@host# **edit** filter *filter-name*

9. Specify the interface-shared filter.

   [edit dynamic-profiles *profile-name* firewall family *family-name* filter *filter-name*]
   user@host# **set interface-shared**

In the following example using an interface-shared filter, you configure a dynamic profile that is used to implement agent-circuit-identifier VLAN household filtering. If **$junos-input-filter** is FILTER1 and **$junos-interface-set-name** is ACI1, then a filter with the name FILTER1-ACI1-in is created and attached to the demux0 unit. When a subsequent login from the same household occurs, it is in the same VLAN. If **$junos-input-filter** is also FILTER1, the next demux0 interface also has the FILTER1-ACI1-in filter attached. A low value precedence was used with the interface-shared filter. If you want to have the interface-shared filter applied first, then you must give a higher precedence to any other filters that are attached to the same interfaces.

```
[edit]
  dynamic-profile {
  client-profile {
    interfaces {
      demux0 {
        unit $junos-interface-unit {
          family inet {
            filter {
```

 Copyright © 2013, Juniper Networks, Inc.

```
                    input $junos-input-filter shared-name $junos-interface-set-name precedence
                      10;
                  }
                }
              }
            }
          }
        }
        }
        firewall {
          family inet {
            filter FILTER1 {
              interface-shared;
              term... # the filter's terms
            }
          }
        }
```

Related Documentation

- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type on page 1169
- Dynamically Attaching Filters Using RADIUS Variables on page 1172
- For information about firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*

Example: Configuring Unicast RPF in a Dynamic Profile on MX Series Routers

This example shows how to help defend the router ingress interfaces against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by configuring unicast reverse-path forwarding (RPF) on a customer-edge interface to filter incoming traffic. Unicast RPF verifies the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled. Packets that fail verification are silently discarded unless a fail filter performs some other action on them.

- Requirements on page 1219
- Overview on page 1220
- Configuration on page 1221
- Verification on page 1224

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.3
- An MX Series 3D Universal Edge router

Before you begin:

- Configure the dynamic profile that you intend to use to apply the RPF check.

  See "Configuring a Basic Dynamic Profile" on page 673.

Overview

Large amounts of unauthorized traffic—such as attempts to flood a network with fake service requests in a denial-of-service (DoS) attack—can consume network resources and deny service to legitimate users. One way to help prevent DoS and distributed denial-of-service (DDoS) attacks is to verify that incoming traffic originates from legitimate network sources.

Unicast RPF helps ensure that a traffic source is legitimate (authorized) by comparing the source address of each packet that arrives on an interface to the forwarding-table entry for its source address. If the router uses the same interface that the packet arrived on to reply to the packet's source, this verifies that the packet originated from an authorized source, and the router forwards the packet. If the router does not use the same interface that the packet arrived on to reply to the packet's source, the packet might have originated from an unauthorized source, and the router discards the packet, or passes it to a fail filter.

The fail filter enables you to set criteria for packets you want to be passed in spite of failing the RPF check, such as DHCP packets, which are dropped by default.

On MX Series routers, you can configure unicast RPF in a dynamic profile to apply the configuration to one or more subscriber interfaces. See *Configuring Unicast RPF* for more information about the behavior and limitations of unicast RPF on MX Series routers.

In this example, you configure the router to protect against potential DoS and DDoS attacks from the Internet perpetrated through IPv4 packets arriving on dynamically created VLAN demux interfaces. The dynamic profile, vlan-demux-prof, establishes that VLAN demux interfaces are automatically created for subscribers. Unicast RPF is enabled on the dynamic interfaces by the rpf-check term.

By default, unicast RPF prevents Dynamic Host Configuration Protocol (DHCP) packets from being accepted on interfaces to which it applies. When DHCP packets are discarded, no new subscribers can be created by the dynamic profile. To enable interfaces to accept DHCP packets, you must apply a fail filter that properly sorts through the packets that fail the check and identifies the DHCP packets. In this example, you configure the **allow-dhcp** term in the filter **rpf-pass-dhcp**. This term matches, counts, and accepts IPv4 packets that are destined for the DHCP port and any address. The **default term** drops all other packets that fail the RPF check.

This example does not show all possible configuration choices.

 Copyright © 2013, Juniper Networks, Inc.

Configuration

To enable unicast RPF with a fail filter in a dynamic profile, perform these tasks:

- Configuring the Dynamic Profile to Apply RPF Checking to Dynamic VLAN Demux Interfaces on page 1221
- Configuring the RPF-Check Fail Filter on page 1222

Configuring the Dynamic Profile to Apply RPF Checking to Dynamic VLAN Demux Interfaces

CLI Quick Configuration

To quickly configure the dynamic profile to apply unicast RPF to dynamically created VLAN demux interfaces, 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 dynamic-profiles vlan-demux-prof interfaces demux0
edit unit $junos-interface-unit
set demux-options underlying-interface $junos-interface-ifd-name
set vlan-id $junos-vlan-id
edit family inet
set unnumbered-address lo0.0
set rpf-check fail-filter rpf-pass-dhcp
```

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 unicast RPF on the router:

1. Create a dynamic profile.

   ```
   [edit]
   user@host# edit dynamic-profiles vlan-demux-prof
   ```

2. Specify that the dynamic VLAN profile use the demux interface.

   ```
   [edit dynamic-profiles vlan-demux-prof]
   user@host# edit interfaces demux0
   ```

3. Specify that the dynamic profile applies the demux interface unit value to the dynamic VLANs.

   ```
   [edit dynamic-profiles vlan-demux-prof interfaces demux0]
   user@host# edit unit $junos-interface-unit
   ```

4. Specify the logical underlying interface for the dynamic VLANs.

   ```
   [edit dynamic-profiles vlan-demux-prof interfaces demux0 unit $junos-interface-unit]
   user@host# set demux-options underlying-interface $junos-interface-ifd-name
   ```

5. Configure the variable that results in dynamically created VLAN IDs.

   ```
   [edit dynamic-profiles vlan-demux-prof interfaces demux0 unit $junos-interface-unit]
   user@host# set vlan-id $junos-vlan-id
   ```

6. Configure the IPv4 address family for the demux interfaces.

   ```
   [edit dynamic-profiles vlan-demux-prof interfaces demux0 unit $junos-interface-unit]
   ```

```
user@host# edit family inet
```

7. Configure the unnumbered address for the family.

```
[edit dynamic-profiles vlan-demux-prof interfaces demux0 unit $junos-interface-unit
  family inet]
user@host# set unnumbered-address lo0.0
```

8. Configure unicast RPF and specify the fail filter that is applied to incoming packets that fail the check.

```
[edit dynamic-profiles vlan-demux-prof interfaces demux0 unit $junos-interface-unit
  family inet]
user@host# set fail-filter fail-filter rpf-pass-dhcp
```

Configuring the RPF-Check Fail Filter

CLI Quick Configuration

To quickly configure the unicast RPF-check fail 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 firewall family inet filter rpf-pass-dhcp
edit term allow-dhcp
set from destination-port dhcp
set from destination-address 255.255.255.255/32
set then count rpf-dhcp-traffic
set then accept
up
edit term default
set then discard
```

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 the RPF-check fail filter:

1. Create the fail filter.

```
[edit firewall]
user@host# edit family inet filter rpf-pass-dhcp
```

2. Define the filter term that identifies DHCP packets based on the DHCP destination port, then counts and passes the packets.

```
[edit firewall family inet filter rpf-pass-dhcp]
user@host# edit term allow-dhcp
user@host# set from destination-port dhcp
user@host# set from destination-address 255.255.255.255/32
user@host# set then count rpf-dhcp-traffic
user@host# set then accept
```

3. Define the filter term that drops all other failed packets.

```
[edit firewall filter rpf-pass-dhcp]
user@host# edit term default
user@host# set then discard
```


Copyright © 2013, Juniper Networks, Inc.

Results From configuration mode, confirm the unicast RPF configuration by entering the **show dynamic-profiles** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show dynamic-profiles
vlan-demux-prof {
  interfaces {
    demux0 {
      unit "$junos-interface-unit" {
        vlan-id "$junos-vlan-id";
        demux-options {
          underlying-interface "$junos-interface-ifd-name";
        }
        family inet {
          unnumbered-address lo0.0;
          rpf-check {
            fail-filter rpf-pass-dhcp;
          }
        }
      }
    }
  }
}
```

From configuration mode, confirm the fail filter 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.

```
[edit]
user@host# show firewall
family inet {
  filter rpf-pass-dhcp {
    term allow-dhcp {
      from {
        destination-address {
          255.255.255.255/32;
        }
        destination-port dhcp;
      }
      then {
        count rpf-dhcp-traffic;
        accept;
      }
    }
    term default {
      then {
        discard;
      }
    }
  }
}
```

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

Verification

To confirm that the configuration is correct, perform these tasks:

- Verifying That Unicast RPF Is Enabled on the Router on page 1224

Verifying That Unicast RPF Is Enabled on the Router

Purpose Verify that unicast RPF is enabled.

Action Verify that unicast RPF is enabled by using the **show subscribers extensive** command.

```
user@host> show subscribers extensive
Type: VLAN
  Logical System: default
  Routing Instance: default
  Interface: ae0.1073741824
  Interface type: Dynamic
  Dynamic Profile Name: vlan-demux-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-pass-dhcp
```

Meaning The IPv4 rpf-check Fail Filter Name field displays **rpf-pass-dhcp**, the name of the fail filter applied by the dynamic profile for IPv4 packets failing the RPF check.

Related Documentation

- *Unicast RPF in Dynamic Profiles for Subscriber Interfaces*
- Configuring Unicast RPF and Fail Filters in Dynamic Profiles for Subscriber Interfaces on page 1179
- Configuring a Basic Dynamic Profile on page 673

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 69

Redirecting HTTP Requests Overview

- Redirecting HTTP Requests Overview on page 1225
- Remote HTTP Redirect Server Operation Flow on page 1226
- Local HTTP Redirect Server Operation Flow on page 1227

Redirecting HTTP Requests Overview

HTTP request traffic from subscribers is aggregated from access networks onto a Broadband Remote Access Server (B-RAS) router, where HTTP traffic can be intercepted and redirected to a captive portal. A captive portal provides authentication and authorization services for redirected subscribers before granting access to protected servers outside of a walled garden. A walled garden defines a group of servers where access is provided to subscribers without reauthorization through a captive portal. You can use a captive portal page as the initial page a subscriber sees after logging in to a subscriber session and as a page used to receive and manage HTTP requests to unauthorized Web resources.

The HTTP redirect service implements a data handler and a control handler and registers them with service rules applicable to the HTTP applications. These rules are parsed by the captive-portal-content-delivery process on the routing engine. The data handler applies the rules to HTTP data flows and handles rewriting the IP destination address or sending an HTTP 302 response with a preconfigured redirect URL. In addition, the control handler maintains a connection with the captive-portal-content-delivery process on the routing engine to learn configuration changes, such as the redirect URL and the rewrite IP destination and port pair. To achieve faster performance, the control handler maintains a cache of relevant configured entities, such as URLs on Multiservices DPC.

Packet flow differs depending on the following configurations:

- Walled garden as a service filter—HTTP traffic destined to servers within the walled garden does not flow to Multiservices DPC. However, any HTTP traffic destined outside of the walled garden flows to the Multiservices DPC.
- Walled garden as an HTTP policy term—All HTTP traffic flows to the Multiservices DPC. The HTTP service handler determines whether traffic is allowed to go to a walled garden.
- HTTP request packet—If the flow is destined to servers within the walled garden, no action is taken.

Copyright © 2013, Juniper Networks, Inc.

An HTTP redirect service can be attached to either a static or dynamic interface. For dynamic subscriber management, HTTP services can be attached dynamically at subscriber login or by using a change of authorization (CoA).

Redundant multiservice PIC and DPC support for HTTP redirect distributes captive portal content delivery rules to both PICs to leverage all framework support (for IPv4 only). Data traffic is sent only to the active PIC and rule processing is performed on the active PIC.

Related Documentation

- Configuring a Basic Dynamic Profile on page 673
- Configuring a Dynamic Profile for Various Levels of Services on page 681
- Junos OS Predefined Variables on page 642
- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197

Remote HTTP Redirect Server Operation Flow

You can use the remote HTTP redirect feature in configurations where the redirect server resides outside of the router and on a policy server, such as Session and Resource Control (SRC).

An HTTP redirect remote server that resides in a walled garden behind routers processes HTTP requests redirected to it and responds with a redirect URL to a captive portal. When you use a remote HTTP redirect server, you need to configure an HTTP service rule to rewrite the IP-DA of the incoming HTTP requests on the service router so that the requests reach the remote HTTP redirect server before being redirected to a captive portal.

The following general sequence occurs during access configuration for a remote HTTP redirect server deployment:

1. The subscriber logs in.
2. RADIUS authenticates the subscriber and sends a service activate (IP-DA rewrite), which redirects traffic to the redirect policy server in a walled garden.
3. The subscriber attempts to access the content server.
4. The router first redirects the HTTP traffic to SRC, which redirects it to the captive portal.
5. The captive portal sends an authorization page back to the subscriber.
6. The subscriber enters credentials to obtain authorization.
7. The captive portal verifies the subscriber credentials.
8. The captive portal authorizes the subscriber and notifies SRC.
9. SRC checks the subscriber database and formulates a policy to allow the subscriber access to the content server.

 Copyright © 2013, Juniper Networks, Inc.

10. SRC sends the policy directly to the router or notifies the RADIUS server, which in turn sends a change of authorization (COA) to the router.
11. The router attaches the new policy, overriding the initial IP-DA write.

The subscriber now has access to the content server.

The following example shows a configuration for IP-DA rewrite:

```
[edit services captive-portal-content-delivery]
rule ipda-rewrite {
  match-direction input-output;
  term 1 {
    from {
      applications http {
        destination-port 80;
      }
    }
    then {
      rewrite destination-address 100.20.1.2;
    }
  }
}
```

Related Documentation

- Local HTTP Redirect Server Operation Flow on page 1227

Local HTTP Redirect Server Operation Flow

You can use the local HTTP redirect feature in configurations where the redirect server resides locally on the router.

An HTTP redirect local server that resides locally on a router processes HTTP requests redirected to it and responds with a redirect URL to a captive portal. You can implement the local server as a service within a service set, which provides more scalability and better performance. When you use a local HTTP redirect server, you need to configure an HTTP service rule to redirect HTTP requests to a captive portal within a walled garden.

The following general sequence occurs during access configuration for a local HTTP redirect server deployment:

1. The subscriber logs in.
2. RADIUS authenticates the subscriber and sends a service activate (HTTP redirect), which redirects HTTP traffic to the captive portal in a walled garden.
3. The subscriber attempts to access the content server (HTTP traffic).
4. The subscriber's HTTP traffic is redirected to the captive portal by the router.
5. The captive portal sends an authorization page back to the subscriber.
6. The subscriber enters credentials to obtain authorization.
7. The captive portal verifies the subscriber credentials.
8. The captive portal authorizes the subscriber.

Copyright © 2013, Juniper Networks, Inc.

The subscriber now has access to the content server.

The following example shows a configuration for HTTP redirect:

```
[edit services captive-portal-content-delivery]
rule redirect {
   match-direction input-output
   term 1 {
  from {
    applications junos-http;
  }
  then {
    redirect http://100.20.2.10/index.html; # this is the captive portal page        }
   }
}
```

Related Documentation

- Remote HTTP Redirect Server Operation Flow on page 1226

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 70

Configuring HTTP Redirect

- Configuring HTTP Redirect Services on page 1229
- Verifying HTTP Redirect Requests on page 1232

Configuring HTTP Redirect Services

You can configure a walled garden with services and policies.

To configure the HTTP redirect service:

1. Configure the packet and installation.

```
[edit chassis]
fpc 1 {
  pic 0 {
    adaptive-services {
      service-package {
        extension-provider {
          control-cores 1;
          data-cores 7;
          object-cache-size 1024;
          policy-db-size 64;
          package jservices-cpcd;
          syslog {
            daemon any;
            external any;
          }
        }
      }
    }
  }
}
```

2. Configure the units and assign the VLAN IDs.

```
[edit interfaces]
ge-0/0/1 {
  vlan-tagging;
  unit 1 {
    vlan-id 100;
    family inet {
      address 100.20.1.1/24;
    }
```

```
    }
  }
```

3. Configure the policy options.

```
policy-options {
  prefix-list google {
    74.125.19.0/24;
  }
}
```

4. Configure the service options.

```
firewall {
  family inet {
    service-filter walled {
      term google {
        from {
          destination-prefix-list {
            google;
          }
        }
        then skip;
      }
      term http {
        from {
          destination-port [ 80 8080 443 ];
        }
        then service;
      }
      term skip {
        then skip;
      }
    }
    service-filter fromSRC {
      term SRC {
        from {
          source-address {
            10.1.2.3/32;
          }
          source-port 8800;
        }
        then service;
      }
      term skip {
        then skip;
      }
    }
    service-filter test {
      term t1 {
        from {
          protocol icmp;
        }
        then service;
      }
    }
  }
```

 Copyright © 2013, Juniper Networks, Inc.

```
}
```

5. Configure the captive portal content delivery services.

```
services {
  captive-portal-content-delivery {
    rule test {
      match-direction input;
      term t1 {
        then {
          rewrite;
        }
      }
    }
    profile ipda-rewrite {
      cpcd-rules test;
      ipda-rewrite-options {
        destination-address 10.1.2.3;
        destination-port 8800;
      }
    }
    traceoptions {
      file cpcdd;
      flag all;
    }
  }
  service-set sset1 {
    captive-portal-content-delivery-profile ipda-rewrite;
    interface-service {
      service-interface ms-1/0/0;
    }
  }
  stateful-firewall {
    rule Rule1 {
      match-direction input-output;
      term 1 {
        from {
          applications [ junos-icmp-all junos-dhcp-server junos-tftp junos-http ];
        }
        then {
          accept;
        }
      }
      term 2 {
        from {
          applications SRC;
        }
        then {
          accept;
        }
      }
    }
  }
}
```

6. Configure the applications.

```
applications {
  application SRC {
    protocol tcp;
    destination-port 8800;
  }
}
```

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

Verifying HTTP Redirect Requests

Purpose View information and statistics for the HTTP redirect configuration.

Action

- To display services statistics:

  user@host> **show services captive-portal-content-delivery statistics**

- To display services flows:

  user@host> **show services captive-portal-content-delivery flows**

- To clear services statistics:

  user@host> **clear services captive-portal-content-delivery statistics**

Related Documentation

- *Configuring HTTP Redirect Services*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 71

HTTP Redirect Examples

- Example: Walled Garden as a Service Filter on page 1233
- Example: Walled Garden as an HTTP Service Rule on page 1234
- Example: HTTP Service Within a Service Set on page 1235
- Example: HTTP Service Attached to a Static Interface on page 1235
- Example: HTTP Service Attached to a Dynamic Interface on page 1237
- Example: Configuring Redundant Multiservice on page 1239
- Example: Configuring Destination Address Rewrite for HTTP Redirect on page 1240

Example: Walled Garden as a Service Filter

Service filters are configured under the firewall and are not specific to captive portal content delivery. The following example shows a walled garden with one server, which is the captive portal:

```
[edit firewall family inet]
root@host# show
service-filter walled {
   term 1 {
      from {
         destination-address {
            100.20.2.3/32; ## this is the address of captive portal
         }
         destination-port 80;
      }
      then skip; ## skip service DPC for http traffic
      ## destined to captive portal
   }
}
```

The following example shows a walled garden within a subnet:

```
service-filter walled-net {
   term 2 {
      from {
         destination-prefix-list {
            100.20.2.0/24; ## '100.20.2.0/24' is not defined
         }
      }
      then skip;
```

```
    }
  }
```

The following example shows the configuration of an IPv6 walled garden:

```
[edit services captive-portal-content-delivery]
  rule walled-garden {
  match-direction input-output
  term 1 {
    from {
      destination-address 2001:2002:0:1::/64; ## captival portal resides here
      destination-port 80;
    }
    then {
      accept;
    }
  }
}
```

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

Example: Walled Garden as an HTTP Service Rule

HTTP service rule configuration resides under the services hierarchy and uses the captive portal and content delivery (captive-portal-content-delivery) service. The following example shows a walled garden configured as an HTTP service rule:

```
[edit services captive-portal-content-delivery]
rule walled-garden {
  match-direction input-output
  term 1 {
    from {
      destination-address 100.20.2.3/32; ## captive portal
      destination-port 80;
    }
    then {
      accept;
    }
  }
}
```

When a remote HTTP redirect server is used, you need to configure an HTTP service rule to rewrite the IP-DA of incoming HTTP requests on the service router so that the requests reach the remote HTTP redirect server before being redirected to a captive portal. If the destination port is not specified, the default behavior is determined by the rewrite configuration. If no rewrite configuration is available, the destination port is not rewritten. The following example shows a configuration for IP-DA rewrite:

```
[edit services captive-portal-content-delivery]
rule ipda-rewrite {
  match-direction input-output;
  term 1 {
    from {
      applications junos-http;
```

 Copyright © 2013, Juniper Networks, Inc.

```
        }
        then {
          rewrite destination-address 100.20.2.10; # this is the remote
          # redirect server.
        }
      }
    }
```

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

Example: HTTP Service Within a Service Set

To become part of a service set, you must configure an HTTP service rule under a service set. In the following example, you can use http-service as an option in the service order configuration:

```
[edit services]
service-set http-redirect-walled {
  cpcd-rules walled-garden;
  cpcd-rules redirect;
}
```

You can also put rules in a rule set and then configure the service set as in the following example:

```
[edit services]
service-set http-redirect-walled {
  cpcd-rule-sets redirect-with-walled-garden;
}
```

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

Example: HTTP Service Attached to a Static Interface

The following example shows an HTTP service set attached to a static interface:

```
[edit interfaces ge-1/0/1]
root@hostr# show
unit 0 {
  family inet {
    service {
      input {
        service-set http-redirect-walled;
      }
      output {
        service-set http-redirect-walled;
      }
    }
    address 10.1.3.2/24;
  }
}
```

The following example uses a service filter as a walled garden by defining a rule named redirect, referencing the rule in a profile named http-redirect, configuring a service set named http-redirect that references the http-redirect captive portal content delivery profile, and attaching the http-redirect service set to static interface ge-1/0/1.0.

```
[edit services]
captive-portal-content-delivery {
   rule redirect {
      match-direction input;
      term t1 {
         from {
            destination-address {
               100.0.1.1/32;
            }
         }
         then {
            redirect http://www.google.com;
         }
      }
   }
   profile http-redirect {
      cpcd-rules redirect;
   }
}
service-set http-redirect {
   captive-portal-content-delivery-profile http-redirect;
   interface-service {
      service-interface ms-1/0/0;
   }
}
```

```
[edit interfaces ge-1/0/1]
unit 0 {
   family inet {
      service {
         input {
            service-set http-redirect service-filter walled;
         }
         output {
            service-set http-redirect;
         }
      }
      address 10.1.3.2/24;
   }
}
```

The following example shows an IPv6 static service attachment:

```
[edit interfaces ge-1/0/1]
unit 0 {
   family inet6 {
      service {
         input {
            service-set http-redirect6 service-filter walled6;
         }
         output {
```

 Copyright © 2013, Juniper Networks, Inc.

```
            service-set http-redirect6 service-filter walled6;
          }
        }
        address 2001:2002::1;
      }
    }
```

This example configures the service filter for walled6:

```
firewall {
  family inet6 {
    service-filter walled6 {
      term google {
        from {
          destination-prefix-list {
            google6;
          }
        }
        then skip;
      }
      term http {
        from {
          destination-port [ 80 8080 443 ];
        }
        then service;
      }
      term skip {
        then skip;
      }
    }
  }
}
```

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

Example: HTTP Service Attached to a Dynamic Interface

A dynamic service attachment uses a dynamic profile. In the following dynamic profile example, the name of the service set can be populated dynamically for each subscriber at instantiation time. This dynamic profile encapsulates a service attachment point associated with a statically preprovisioned service set sset-1.

```
dynamic-profiles {
  profile prof-2 { # parameterized service attachment
    interfaces {
      $junos-interface-ifd-name {
        unit $junos-interface-unit {
          family inet {
            service {
              input {
                service-set $junos-service-set service-filter $junos-service-filter;
                post-input-filter $junos-post-input-filter ;
              }
              output {
```

```
                service-set $junos-service-set;
              }
            }
          }
        }
      }
    }
  }
```

To handle scalability more efficiently, in the following example the name of the service set can be populated dynamically for each subscriber at instantiation time.

```
dynamic-profiles {
  profile prof-2 { # parameterized service attachment
    interfaces {
      $junos-interface-ifd-name {
        unit $junos-interface-unit {
          family inet {
            service {
              input {
                service-set $junos-service-set service-filter $junos-service-filter;
                post-input-filter $junos-post-input-filter ;
              }
              output {
                service-set $junos-service-set;
              }
            }
          }
        }
      }
    }
  }
```

The following attaches a service set dynamically at family inet6:

```
dynamic-profiles {
  profile prof-1 {
    interfaces {
      $junos-interface-ifd-name {
        unit $junos-interface-unit {
          family inet6 {
            service {
              input {
                service-set sset-1 service-filter fltr-1;
                post-input-filter pfltr-1 ;
              }
              output {
                service-set sset-1 service-filter fltr-1;
              }
            }
          }
        }
      }
    }
```

 Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
```

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

Example: Configuring Redundant Multiservice

- Requirements on page 1239
- Overview on page 1239
- Configuration on page 1239
- Verification on page 1240

Requirements

- Multiservices DPC PIC

Overview

This procedure shows how to configure redundant multiservice support.

Configuration

Example: Configuring Redundant Multiservice for IPv4

Step-by-Step Procedure

1. Configure the interface:

   ```
   [edit interfaces]
   user@host# set interface rms0
   ```

2. Configure the redundant multiservice service set:

   ```
   [edit services]
   user@host# set service-interface rms0
   ```

3. Configure the redundant multiservice service set attachment:

   ```
   [edit interfaces]
   user@host# set ge-1/0/0 unit 100
   ```

Results

Confirm the configuration by entering the **show redundancy-options** configuration command.

```
show redundancy-options
  redundancy-options {
    primary ms-2/1/0;
    secondary ms-3/1/0;
    hot-standby;
  }
    unit 0 {
      family inet;
    }
```

Confirm the service set configuration by entering the **show captive-portal-content-delivery-profile** configuration command.

```
show captive-portal-content-delivery-profile httpRedirect
  interface-service {
    service-interface rms0;
  }
```

Confirm the service set attachment by entering the **show show vlan-id** configuration command.

```
show vlan-id 100
  family inet {
    service {
      input {
        service-set sset10 service-filter walled;
      }
      output {
        service-set sset10;
      }
    }
    address 192.1.4.1/24;
  }
```

Verification

Displaying Redundant Multiservice Configuration

Purpose Verify the redundant multiservice configuration.

Action user@host> **show interfaces redundancy detail**

Related Documentation
- *Failover of the Control Service PICs*

Example: Configuring Destination Address Rewrite for HTTP Redirect

- Requirements on page 1240
- Overview on page 1240
- Configuration on page 1241
- Verification on page 1242

Requirements

- Multiservices DPC PIC

Overview

This procedure shows how to configure an DA rewrite rule. The destination port is not specified and the default behavior is determined by the rewrite configuration. If no rewrite configuration is available, the destination port is not rewritten.

 Copyright © 2013, Juniper Networks, Inc.

Configuration

Example: Configuring a Rewrite Rule

Step-by-Step Procedure

1. Configure the service rule:

```
[edit services captive-portal-content-delivery]
user@host# set rule da-rewrite
```

2. Specify the term name:

```
[edit services captive-portal-content-delivery da-rewrite]
user@host# set term t1
```

3. Specify the match conditions for the term:

```
[edit services captive-portal-content-delivery da-rewrite inet-filter term t1]
user@host# set from applications junos-http
```

4. Specify the actions to take if the packet matches all the conditions in that term:

```
[edit services captive-portal-content-delivery da-rewrite inet-filter term t1]
user@host# set then rewrite desitnation-address 2001:2002::1;
```

Results

Confirm the configuration by entering the **show services** configuration command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit services captive-portal-content-delivery]
rule da-rewrite {
   match-direction input-output
   term 1 {
      from {
         applications junos-http;
      }
      then {
         rewrite destination-address 2001:2002::1; # this is the remote redirect server.
      }
   }
}
```

The following example shows the configuration for an IPv6-DA rewrite service rule. Because the destination port is not specified, the default behavior is determined by the rewrite configuration. If no rewrite configuration is available, the destination port is not rewritten.

```
[edit services captive-portao-content-delivery]
rule ipv6da-rewrite {
   match-direction input-output
   term 1 {
      from {
         applications junos-http;
      }
      then {
         rewrite destination-address 2001:2002::1; # this is the remote
         # redirect server.
      }
```

```
    }
}
```

Verification

Displaying HTTP Redirect configuration

Purpose Verify the HTTP requests are redirected to the server.

Action user@host> **show services detail**

Related Documentation

- *Failover of the Control Service PICs*

Copyright © 2013, Juniper Networks, Inc.

PART 16

Subscriber Secure Policy Traffic Mirroring

- Subscriber Secure Policy Traffic Mirroring Overview on page 1245
- Overview of Subscriber Secure Policy Using RADIUS on page 1247
- Configuring RADIUS-Initiated Subscriber Secure Policy Traffic Mirroring on page 1261
- Overview of Subscriber Secure Policy Using DTCP on page 1267
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring on page 1281
- DTCP Messages Used for DTCP-Initiated Subscriber Secure Policy Mirroring on page 1291
- Reporting Intercept-Related Events for Subscriber Secure Policy on page 1305

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 72

Subscriber Secure Policy Traffic Mirroring Overview

- Subscriber Secure Policy Overview on page 1245
- Subscriber Secure Policy Licensing Requirements on page 1246

Subscriber Secure Policy Overview

Subscriber secure policy enables you to mirror traffic on a per-subscriber basis. You can mirror the content of subscriber traffic as well as monitor events related to the subscriber session that is being mirrored.

Subscriber secure policy mirroring can be based on information provided by either RADIUS or Dynamic Tasking Control Protocol (DTCP), and can mirror both IPv4 and IPv6 traffic. Configuration of subscriber secure policy mirroring is independent of the actual mirroring session—you can configure the mirroring parameters at any time. Also, you can use a single RADIUS or DTCP server to provision mirroring operations on multiple routers in a service provider's network. To provide security, the ability to configure, access, and view the subscriber secure policy components and configuration is restricted to authorized users.

After subscriber secure policy is triggered, both the subscriber incoming and outgoing traffic are mirrored. The original traffic is sent to its intended destination and the mirrored traffic is sent to a mediation device for analysis. The actual mirroring operation is transparent to subscribers whose traffic is being mirrored. A special UDP/IP header is prepended to each mirrored packet sent to the mediation device. The mediation device uses the header to differentiate multiple mirrored streams that arrive from different sources.

Subscriber Secure Policy for Subscribers on VLANs

Interface-based subscriber secure policy is supported on dynamic, authenticated VLAN interfaces and VLAN demux interfaces. When you enable subscriber secure policy for these interfaces, traffic for all configured families (inet, inet6) including Layer 2 and Layer 3 control traffic is mirrored. The mirrored packets include Layer 2 encapsulations.

Copyright © 2013, Juniper Networks, Inc.

Traffic Filtering For DTCP-Initiated Subscriber Secure Policy Mirrored Traffic

You can filter mirrored traffic before it is sent to a mediation device. With this feature, service providers can reduce the volume of traffic sent to a mediation device. For some types of traffic, such as IPTV or video on demand, you do not need to mirror the entire content of the traffic because the content may already be known or controlled by the service provider.

Mirroring-Related Event Reporting

Subscriber secure policy also supports the use of SNMPv3 traps to report events related to the mirroring operation to an external device. Types of information sent in traps include identifying information for subscribers, such as username or IP address, and subscriber session events, such as login or logout events or mirroring session activation or deactivation. The traps map to messages defined in the *Lawfully Authorized Electronic Surveillance (LAES) for IP Network Access, American National Standard for Telecommunications*.

Related Documentation

- RADIUS-Initiated Subscriber Secure Policy Overview on page 1247
- DTCP-Initiated Subscriber Secure Policy Overview on page 1267
- Intercept-Related Events Transmitted to the Mediation Device on page 1305

Subscriber Secure Policy Licensing Requirements

To enable and use subscriber secure policy, you must install and properly configure the Subscriber Secure Policy license.

Related Documentation

- For information about installing and managing Junos licenses, see “Installing and Managing Junos OS Licenses” in the *Installation and Upgrade Guide*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 73

Overview of Subscriber Secure Policy Using RADIUS

- RADIUS-Initiated Subscriber Secure Policy Overview on page 1247
- Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS on page 1248
- RADIUS-Initiated Traffic Mirroring Interfaces on page 1250
- RADIUS-Initiated Traffic Mirroring Process at Subscriber Login on page 1251
- RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers on page 1252
- Subscriber Secure Policy Support for IPv4 Multicast Traffic on page 1253
- RADIUS Attributes Used for Subscriber Secure Policy on page 1254
- Using the Packet Header to Track Subscribers on the Mediation Device on page 1255
- Subscriber Secure Policy and L2TP LAC Subscribers on page 1259
- Subscriber Secure Policy and L2TP LNS Subscribers on page 1259

RADIUS-Initiated Subscriber Secure Policy Overview

RADIUS-initiated mirroring creates secure policies based on RADIUS VSAs and uses RADIUS attributes to identify the subscriber whose traffic is to be mirrored. Mirroring is initiated without regard to the subscriber location, router, interface, or type of traffic.

The mirroring operation can be initiated by RADIUS messages as follows:

- Subscriber login—Mirroring starts when the subscriber logs in and the router receives the trigger in a RADIUS Access-Accept message. Using triggers in RADIUS Access-Accept messages enables you to mirror per-subscriber traffic without regard to how often the subscriber logs in or out, or which router or interface the subscriber uses.
- In-session—Mirroring starts when the router receives the trigger in a RADIUS change of authorization request (CoA-Request) message. Using triggers in CoA-Request messages enables you to immediately mirror traffic of a subscriber who is already logged in.

Related Documentation

- Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS on page 1248
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

Copyright © 2013, Juniper Networks, Inc.

Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS

Figure 38 on page 1248 shows the architecture of the RADIUS-initiated subscriber secure policy mirroring environment.

Figure 38: RADIUS-Initiated Subscriber Secure Policy Architecture

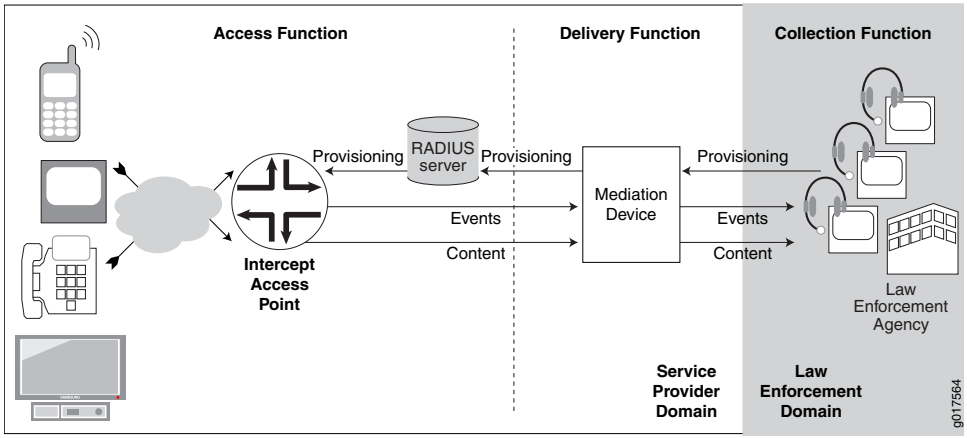

Table 109 on page 1248 describes the functions and components of a RADIUS-initiated subscriber secure policy traffic mirroring environment.

Table 109: RADIUS-Initiated Subscriber Secure Policy Functions and Components

Function or Component	Description
Collection function	The collection function is responsible for collecting intercepted content and identifying information from the delivery function. The collection function is the responsibility of the law enforcement agency (LEA).
Delivery function	The delivery function delivers information that it receives from the access function to the collection function. The delivery function is performed by the mediation device.
Access function	The access function has access to the intercept target's traffic content and intercept-related events. It is responsible for collecting this information and sending it to the delivery function. The access function is the responsibility of intercept access points (IAPs).

Copyright © 2013, Juniper Networks, Inc.

Table 109: RADIUS-Initiated Subscriber Secure Policy Functions and Components *(continued)*

Function or Component	Description
Events	Intercept-related events, such as login or logout events or mirroring session activation or deactivation. The router sends the events to the mediation device in SNMP traps.
LEA	Law enforcement agency. The LEA provides intercept targets to the service provider who provisions the mediation device.
Mediation device	The mediation device receives provisioning information from the LEA, and it uses the information to send provisioning information to the RADIUS server. The mediation device also receives intercept-related events and intercepted content from the router, and delivers the events and intercepted content to the LEA.
RADIUS server	The RADIUS server receives provisioning information from the mediation device. It identifies subscribers whose traffic is to be mirrored, and triggers mirroring sessions on the IAP (the router) by including mirroring-related RADIUS attributes and VSAs in Access-Accept or CoA-Request messages that it sends to the IAP.
IAP	Intercept access point. In a subscriber access network the Juniper Networks router is the IAP. Using subscriber secure policies, the IAP intercepts traffic to and from the subscriber whose traffic is being mirrored. It encapsulates the intercepted content in a packet header and delivers it to the mediation device, while also sending the content to the intended destination. The IAP also sends intercept-related events to the mediation device using SNMP traps.

Related Documentation

- RADIUS-Initiated Subscriber Secure Policy Overview on page 1247
- RADIUS-Initiated Traffic Mirroring Interfaces on page 1250
- RADIUS-Initiated Traffic Mirroring Process at Subscriber Login on page 1251
- RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers on page 1252

Copyright © 2013, Juniper Networks, Inc.

RADIUS-Initiated Traffic Mirroring Interfaces

Figure 39 on page 1250 shows the interfaces involved in RADIUS-initiated secure subscriber policy traffic mirroring.

Figure 39: RADIUS-Initiated Traffic Mirroring Interfaces

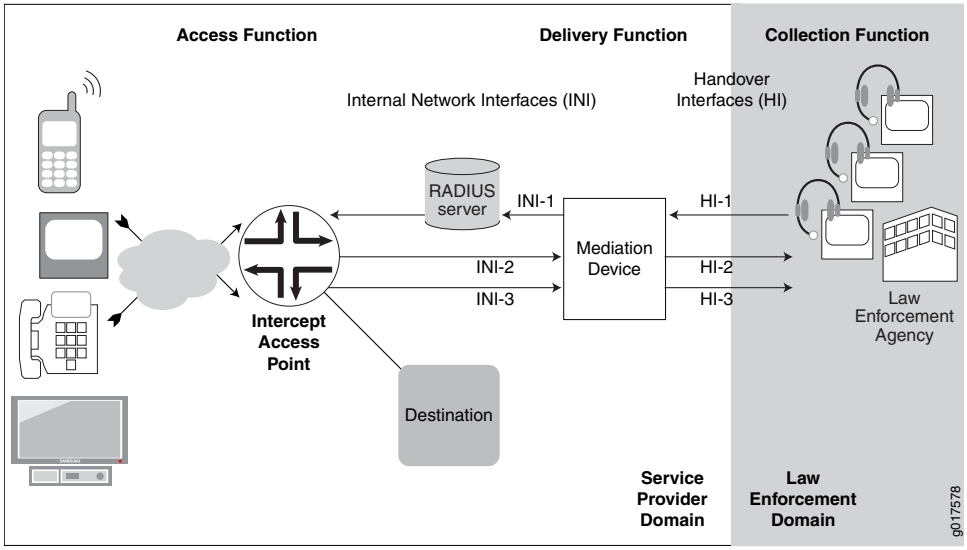

Table 110 on page 1250 describes the interfaces involved in RADIUS-initiated secure subscriber policy traffic mirroring.

Table 110: RADIUS-Initiated Traffic Mirroring Interfaces

Interface	Description
HI-1	Handover Interface 1—Administrative interface between the LEA and the service provider mediation device. The LEA sends provisioning information to the mediation device on this interface.
HI-2	Handover Interface 2—Intercept-related information interface between the LEA and the mediation device that is used to deliver intercept-related events to the LEA. These events can be subscriber session events such as login, logout, and authentication.
HI-3	Handover Interface 3—Intercepted content Interface between the mediation device and LEA that is used to deliver intercepted content to the LEA.
INI-1	Internal network Interface 1—Interface used to send intercept provisioning information from the mediation device to the RADIUS server.
INI-2	Internal network interface 2—Interface used to send intercept-related events from the router to the mediation device. This information is sent in SNMP traps.
INI-3	Internal network interface 3—Interface used to send intercepted content from the router to the mediation device.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS on page 1248
- RADIUS-Initiated Traffic Mirroring Process at Subscriber Login on page 1251
- RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers on page 1252

RADIUS-Initiated Traffic Mirroring Process at Subscriber Login

Figure 40 on page 1251 shows the process for a RADIUS-initiated subscriber mirroring operation that is initiated when the mirrored subscriber logs in.

Figure 40: RADIUS-Initiated Subscriber Secure Policy Model at Login

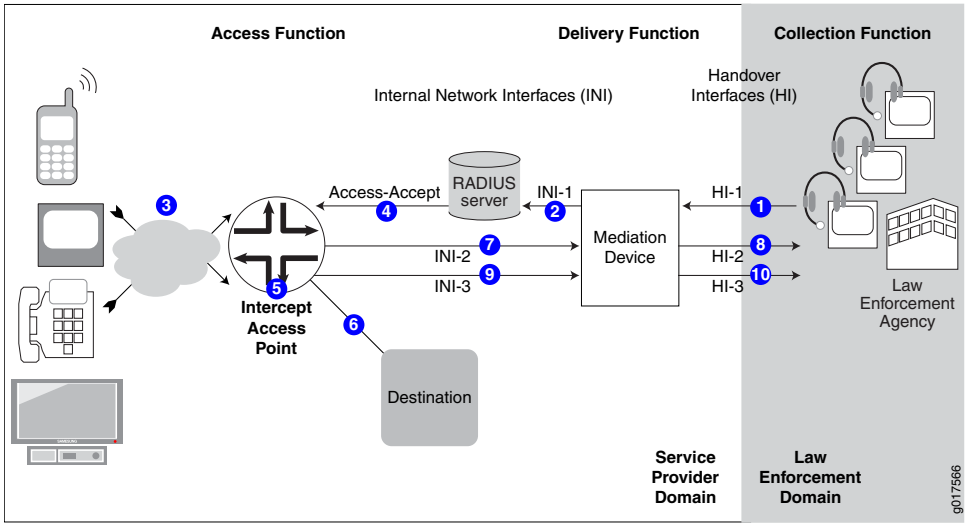

1—The LEA sends provisioning information for a subscriber whose traffic is to be mirrored over the HI-1 interface to the mediation device.	6—The IAP sends the original subscriber traffic to its intended destination.
2—The mediation device sends the provisioning information over the INI-1 interface to the RADIUS server.	7—As subscriber-related events occur, the IAP sends the events in SNMP traps over the INI-2 interface to the mediation device.
3—The subscriber logs in, requesting authentication by the RADIUS server.	8—The mediation device provides the events over the HI-2 interface to the LEA.
4—The RADIUS server authenticates the subscriber and sends an Access-Accept message containing mirroring-related RADIUS attributes in Juniper Networks VSAs to the IAP (the router).	9—The IAP encapsulates the mirrored content in a packet header and sends it over the INI-3 interface to the mediation device. The IAP uses the destination IP address of the mediation device that it received in the Access-Accept messaged from the RADIUS server.
5—The IAP creates a subscriber secure policy based on the mirroring VSAs and begins mirroring the subscriber's traffic.	10—The mediation device sends mirrored content over the HI-3 interface to the LEA.

Related Documentation

- Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS on page 1248
- RADIUS-Initiated Traffic Mirroring Interfaces on page 1250

- RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers on page 1252
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers

Figure 41 on page 1252 shows the process for a RADIUS-initiated subscriber mirroring operation that is initiated after the subscriber has logged in.

Figure 41: RADIUS-Initiated Subscriber Secure Policy Model After Login

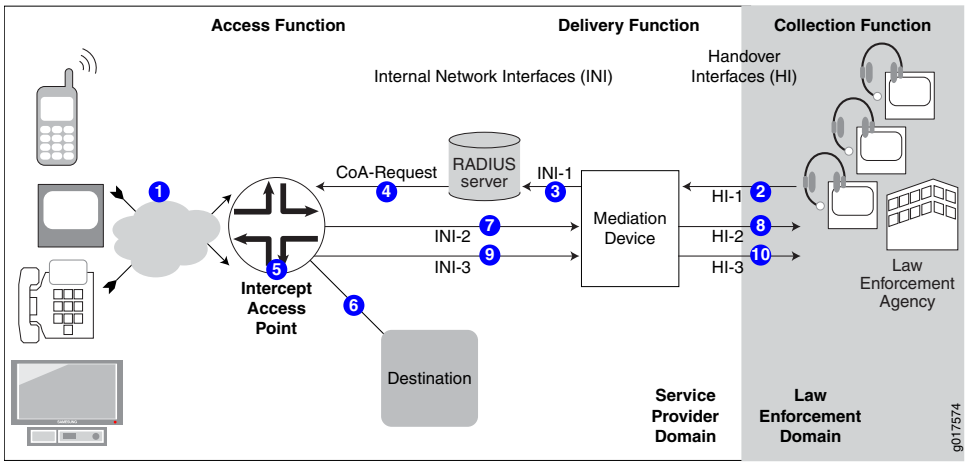

1—The subscriber logs in, requesting authentication by the RADIUS server. The RADIUS server authenticates the subscriber (no mirroring activity occurs).	6—The IAP sends the original subscriber traffic to its intended destination.
2—The LEA sends provisioning information for a subscriber whose traffic is to be mirrored over the HI-1 interface to the mediation device.	7—As subscriber-related events occur, the IAP sends the events in SNMP traps over the INI-2 interface to the mediation device.
3—The mediation device sends the provisioning information over the INI-1 interface to the RADIUS server.	8—The mediation device provides events over the HI-2 interface to the LEA.
4—The RADIUS server sends a CoA message containing the mirroring-related RADIUS attributes and VSAs to the IAP (the router).	9—The IAP encapsulates the mirrored subscriber content in a packet header and sends it to the mediation device over the INI-3 interface. The IAP uses the destination IP address that it received in the Access-Accept messaged from the RADIUS server.
5—The RADIUS CoA message initiates the mirroring operation. The IAP creates the subscriber secure policy based on the mirroring VSAs and immediately begins mirroring subscriber traffic.	10—The mediation device sends mirrored content over the HI-3 interface to the LEA.

Related Documentation

- Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS on page 1248
- RADIUS-Initiated Traffic Mirroring Interfaces on page 1250

Copyright © 2013, Juniper Networks, Inc.

- RADIUS-Initiated Traffic Mirroring Process at Subscriber Login on page 1251
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

Subscriber Secure Policy Support for IPv4 Multicast Traffic

IP multicast traffic is used for applications such as audio or video streaming, IPTV, video conferencing, or online gaming. Multicast traffic is sent to multiple subscribers who have joined a multicast group.

Secure subscriber policy allows for the mirroring of IPv4 multicast traffic sent to a specific subscriber. If multiple subscribers whose traffic requires mirroring join the same multicast session, the subscriber secure policy feature mirrors each subscriber's traffic and forwards it separately to the mediation device with the proper prepended header.

Mirroring of multicast traffic is supported only for subscribers in the default logical system.

You can enable and disable the mirroring of multicast traffic on a per-chassis basis. You cannot enable or disable it on a per-subscriber basis.

Triggering the Mirroring of IPv4 Multicast Traffic

Multicast traffic being sent towards a subscriber does not contain much of the identifying information used to trigger mirroring of a subscriber's unicast traffic. For example, the multicast packet contains the multicast group address in the destination address of the packet instead of the subscriber's IP address. It also does not contain the user name or MAC address of the subscriber, and does not include information obtained by RADIUS or DHCP. Therefore, methods of identifying multicast traffic that is received by a subscriber are not the same as methods of identifying a subscriber's unicast traffic or multicast traffic that is sent by a subscriber.

To join a multicast group, a subscriber sends an IGMP join request, and it receives a reply. The reply contains the multicast groups to which the subscriber is registered. Triggering the mirroring of multicast traffic is based on the sending of the IGMP join request and the information in the IGMP reply. If the subscriber's unicast traffic is already being mirrored either through DTCP-initiated or RADIUS-initiated traffic mirroring, and the subscriber sends an IGMP join request, mirroring of multicast traffic sent to the subscriber is initiated. The traffic being mirrored is based on the groups contained in the IGMP reply.

Related Documentation

- Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic on page 1266

RADIUS Attributes Used for Subscriber Secure Policy

Table 111 on page 1254 lists the RADIUS VSAs that are associated with subscriber secure policy. If these VSAs are present in the RADIUS Access-Accept message for a subscriber, the action specified in the LI-Action attribute takes effect.

Mirroring VSAs that the RADIUS server sends to the router are salt-encrypted. Salt encryption is a random string of data used to modify a password hash.

Table 111: RADIUS-Based Mirroring Attributes

Attribute Number	Attribute Name	Description	Value
[26-58]	LI-Action	Traffic mirroring action	Salt-encrypted integer • 0 = stop mirroring • 1 = start mirroring • 2 = no action
[26-59]	Med-Dev-Handle	Identifier that associates mirrored traffic with a specific subscriber Med-Dev-Handle includes: • Intercept-Identifier • Acct-Session-ID (optional)	Salt-encrypted string
[26-60]	Med-Ip-Address	IP address of mediation device to which mirrored traffic is forwarded	Salt-encrypted IP address
[26-61]	Med-Port-Number	UDP port in the mediation device to which mirrored traffic is forwarded	Salt-encrypted integer

NOTE: 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.

If a subscriber is already logged in, Table 112 on page 1255 lists the RADIUS attributes that can be present in RADIUS CoA messages to identify the subscriber whose traffic is to have a mirroring action applied (activation or deactivation).

Copyright © 2013, Juniper Networks, Inc.

Table 112: RADIUS Attributes Used in CoA Messages to Identify Subscribers for Traffic Mirroring

Attribute Number	Attribute Name
[1]	User-Name
[44]	Acct-Session-ID

Triggering Subscriber Secure Policy for Subscribers on Dynamic Authenticated VLANs

BEST PRACTICE: When you have DHCPv4/DHCPv6 subscribers over VLANs, two sessions are created for each subscriber—one for the Layer 2 VLAN, and one for DHCP. In this case, we recommend that you use one trigger that matches both the DHCP and the VLAN session.

If authentication is performed on both the VLAN session and the DHCP session, we recommend that you use a separate, unique username for the VLAN and DHCP sessions to allow RADIUS to distinguish on which of the sessions to trigger subscriber secure policy traffic mirroring. Otherwise, when the DHCP session is authenticated and activated, traffic mirroring fails.

Related Documentation

- RADIUS-Initiated Subscriber Secure Policy Overview on page 1247
- Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS on page 1248

Using the Packet Header to Track Subscribers on the Mediation Device

When the router sends mirrored traffic to the mediation device, it encapsulates it in a packet header. Figure 42 on page 1256 is the mirrored packet header and payload that the router sends to the mediation device.

Figure 42: Mirrored Packet Header and Payload

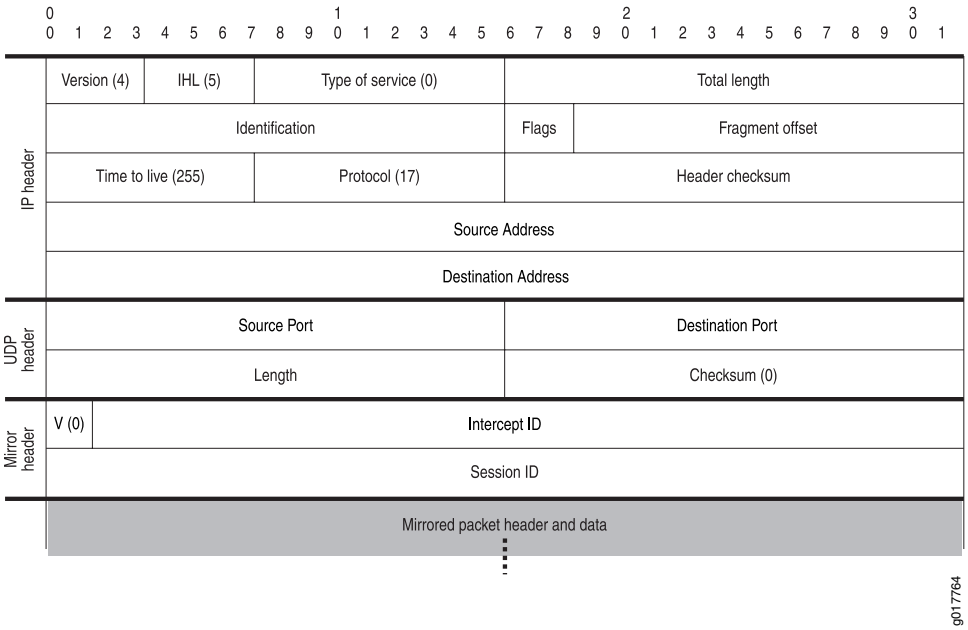

Table 113 on page 1256 describes the fields in the packet header of mirrored packets.

Table 113: Mirrored Packet Header and Payload Field Descriptions

Field	Value	Length (Bits)
IP Header		
Version	4	4
IHL	5	4
Type of Service	0	8
Total Length	Dynamically computed	16
Identification	Dynamically computed	16
Flags	Dynamically computed	3
Fragment Offset	Dynamically computed	13
Time to Live	255	8
Protocol	17	8
Header Checksum	Dynamically computed	16

Copyright © 2013, Juniper Networks, Inc.

Table 113: Mirrored Packet Header and Payload Field Descriptions *(continued)*

Field	Value	Length (Bits)
Source Address	IP address of the router interface that sends mirrored traffic to the mediation device	32
Destination Address	IP address of the mediation device to which mirrored traffic is forwarded (VSA 26-60)	32
UDP Header		
Source Port	UDP port number on the router from which mirrored traffic is sent to the mediation device	16
Destination Port	UDP port on the mediation device to which mirrored traffic is forwarded (VSA 26-61)	16
Length	Dynamically computed	16
Checksum	0	16
Mirror Header		
V (mirror header value)	0	2
Intercept ID	See "Format of the Mirror Header Values Used to Track Subscribers and Subscriber Sessions" on page 1258 for details	30
Session-ID	See "Format of the Mirror Header Values Used to Track Subscribers and Subscriber Sessions" on page 1258 for details	32

Copyright © 2013, Juniper Networks, Inc.

Format of the Mirror Header Values Used to Track Subscribers and Subscriber Sessions

The packet header includes mirror header attributes that the mediation device can use to track subscribers and subscriber sessions. The router creates values for these attributes based on information that it receives from RADIUS. There are three mirror header attributes in the packet header:

- V (mirror header value)—Used by the router to specify how the values of the Session ID and Intercept ID are determined. The value received from RADIUS can be a 0 or a 1. However, the value is always 0 in the packet header sent to the mediation device.
- Session ID—Used by the mediation device to identify the session of the mirrored subscriber. The value is assigned to a subscriber session by the Junos OS. The Session ID changes with each new session for a subscriber.
- Intercept ID—Used along with the Session ID by the mediation device to track a subscriber across multiple login and logout events. The value is assigned to a subscriber whose traffic is being intercepted. The Intercept ID is constant; it does not change as a subscriber logs in and logs out of sessions.

The values of the Intercept ID and the Session ID are determined by the value that the router receives in VSA 26-59. VSA 26-59 is declared as a hexadecimal string that can be either 4 bytes or 8 bytes long. The mirror header value specifies whether a 4-byte value or an 8-byte value is used to form the Intercept ID and the Session ID.

4-Byte Format

The 4-byte format allows you to manually specify the Intercept ID. The Session ID value is automatically created based on the least significant 32 bits of the Acct-Session-ID (RADIUS attribute 44).

To use the 4-byte format of VSA 26-59, you configure the first two most significant bits of the VSA to a value of 1, which indicates a single word in the VSA. The remaining 30 bits of the word form the Intercept ID value.

For example, a value of 40000010 for VSA 26-59 configures the following fields in the mirror header, as shown in Figure 43 on page 1258:

- V = 1
- Intercept ID = 0x10

Figure 43: 4-Byte Format of VSA 26-59

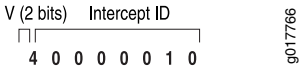

8-Byte Format

The 8-byte format of VSA 26-59 enables you to manually specify the both the Session-ID value and the Intercept ID value.

Copyright © 2013, Juniper Networks, Inc.

To use the 8-byte format, you configure the first two most significant bits of the first word of the VSA to a value of 0, which indicates two words in the VSA. The remaining 30 bits of the first word form the Intercept ID value, and the second word is the Session-ID field. You cannot change the order of these two words.

For example, a value of 0000030000000090 in VSA 26-59 configures the following fields in the mirror header, as shown in Figure 44 on page 1259:

- V = 0
- Intercept-ID = 0x300
- Session-ID = 0x90

Figure 44: 8-Byte Format of VSA 26-59

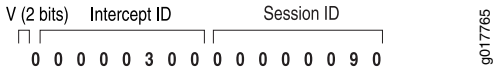

Related Documentation

- RADIUS-Initiated Subscriber Secure Policy Overview on page 1247
- Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS on page 1248

Subscriber Secure Policy and L2TP LAC Subscribers

RADIUS-initiated per-subscriber traffic mirroring can be applied to subscribers whose traffic is tunneled with L2TP. Both subscriber ingress traffic (from the subscriber into the tunnel) and subscriber egress traffic (from the tunnel to the subscriber) are mirrored at the subscriber-facing ingress interface on the LAC. The ingress traffic is mirrored after PPPoE decapsulation and before L2TP encapsulation. The egress traffic is mirrored after L2TP decapsulation. The mirrored packet includes the complete HDLC frame sent to the LNS rather than only the IP datagram.

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- RADIUS Attributes Used for Subscriber Secure Policy on page 1254

Subscriber Secure Policy and L2TP LNS Subscribers

Dynamic Tasking Control Protocol (DTCP)-initiated and RADIUS-initiated per-subscriber traffic mirroring can be applied to Point-to-Point Protocol (PPP) subscribers whose traffic is tunneled with Layer 2 Tunneling Protocol (L2TP). At the L2TP network server (LNS), both subscriber ingress traffic (from the L2TP access concentrator, or LAC, to the LNS) and subscriber egress traffic (from the LNS to the LAC) are mirrored at the inline services (si) interface corresponding to the subscriber. Ingress traffic is mirrored after decapsulation of L2TP, HDLC, and PPP headers. The egress traffic is mirrored before the IP datagram is encapsulated. The mirrored traffic contains only the IP datagram belonging to the subscriber.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- RADIUS Attributes Used for Subscriber Secure Policy on page 1254

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 74

Configuring RADIUS-Initiated Subscriber Secure Policy Traffic Mirroring

- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Guidelines for Configuring Subscriber Secure Policy Mirroring on page 1262
- Configuring Tunnel Interfaces for Subscriber Secure Policy Mirroring on page 1263
- Configuring Support for Subscriber Secure Policy Mirroring on page 1264
- Configuring RADIUS Server Support for Subscriber Secure Policy Mirroring on page 1265
- Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic on page 1266
- Terminating RADIUS-Initiated Subscriber Traffic Mirroring on page 1266

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 tunnel interfaces (vt interfaces) that are used to send mirrored content to the mediation device.

   See "Configuring Tunnel Interfaces for Subscriber Secure Policy Mirroring" on page 1263.

2. 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" on page 1264.

3. 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" on page 1265.

4. 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" on page 1254 for descriptions of the supported attributes used in RADIUS Accept-Accept and CoA messages.
   - The mediation device must be configured to accept the mirrored content.

5. (Optional) Enable the mirroring of IPv4 multicast traffic on the router.

   See "Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic" on page 1266.

6. (Optional) Configure SNMPv3 trap support to report mirroring-related events to the mediation device.

   See "Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring" on page 1307.

To terminate an active subscriber mirroring session at any time.

See "Terminating RADIUS-Initiated Subscriber Traffic Mirroring" on page 1266.

Related Documentation

- RADIUS Attributes Used for Subscriber Secure Policy on page 1254
- Guidelines for Configuring Subscriber Secure Policy Mirroring on page 1262
- Intercept-Related Events Transmitted to the Mediation Device on page 1305
- Terminating RADIUS-Initiated Subscriber Traffic Mirroring on page 1266

Guidelines for Configuring Subscriber Secure Policy Mirroring

The subscriber secure policy service uses the radius-flow-tap service infrastructure.

When configuring subscriber secure policy mirroring, consider the following guidelines regarding the relationship between subscriber secure policy service and the radius-flow-tap service:

- The radius-flow-tap service **[edit services radius-flow-tap]** and the flow-tap service **[edit services flow-tap]** cannot run simultaneously on the router. Therefore, flow-tap and subscriber secure policy mirroring cannot run simultaneously on the same router.
- You can configure one instance of the radius-flow-tap service on the router. Subscriber secure policy RADIUS-initiated mirroring and DTCP-initiated mirroring use the radius-flow-tap service.
- If you delete the radius-flow-tap service all subscriber secure policy mirroring stops.

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281
- Configuring Support for Subscriber Secure Policy Mirroring on page 1264

Configuring Tunnel Interfaces for Subscriber Secure Policy Mirroring

The router, acting as the IAP, uses tunnel interfaces (vt interfaces) to send mirrored traffic to the mediation device. The IAP equally distributes the mirrored traffic across the available tunnel interfaces.

Because the MX Series 3D Universal Edge Routers do not support Tunnel Services PICs, you create a pool tunnel interfaces on MX Series routers at the **[edit chassis]** hierarchy level.

You can configure up to 2048 mirrored subscriber sessions per chassis.

To configure a pool of tunnel interfaces for use by subscriber secure policy mirroring:

1. Access the chassis configuration, and specify the slot number of the DPC, MPC, or MIC.
   - On the MX80 router, the range is 0 through 1.
   - On other MX Series routers, if two System Control Boards (SCBs), are installed, the range is 0 through 11. If three SCBs are installed, the range is 0 through 5 and 7 through 11.

     ```
     [edit chassis]
     user@host# edit fpc 1
     ```
2. Configure the PIC number of the FPC.
   - On MX80 routers, if the FPC is 0, the PIC number can only be 0. If the FPC is 1, the PIC range is 0 through 3.
   - For all other MX Series routers, the range is 0 through 3.

     ```
     [edit chassis fpc 1]
     user@host# edit pic 1
     ```
3. Specify that the FPC and PIC are to be used for tunnel interfaces.

   ```
   [edit chassis fpc 1 pic 1]
   user@host# edit tunnel-services
   ```
4. Specify the amount of bandwidth to reserve for tunnel traffic on each Packet Forwarding Engine.
   - 1g indicates that 1 Gbps of bandwidth is reserved for tunnel traffic.
   - 10g indicates that 10 Gbps of bandwidth is reserved for tunnel traffic.

If you specify a bandwidth that is not compatible, tunnel services are not activated. For example, you cannot specify a bandwidth of 1 Gbps for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC.

```
[edit chassis fpc 1 pic 1 tunnel-services]
user@host#
user@host# set bandwidth 1g
```

5. Configure the tunnel interfaces, including the family.

   To configure subscriber secure policy mirroring for IPv6 traffic, configure the tunnel interfaces for both the **inet** and **inet6** families.

```
[edit interfaces]
user@host# set vt-1/1/0 unit 0 family inet
user@host# set vt-1/1/0 unit 0 family inet6
```

Related Documentation

- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281
- Subscriber Secure Policy and L2TP LNS Subscribers on page 1259

Configuring Support for Subscriber Secure Policy Mirroring

Subscriber secure policy runs on the radius-flow-tap service. This topic describes the steps to configure radius-flow-tap support for RADIUS-initiated and DTCP-initiated subscriber secure policy mirroring.

To configure the radius-flow-tap service to support subscriber secure policy mirroring:

1. Configure the flow-tap service used for subscriber secure policy mirroring.

```
[edit services]
user@host# edit radius-flow-tap
```

2. Assign the tunnel interfaces that the radius-flow-tap service uses.

```
[edit services radius-flow-tap]
user@host# set interfaces vt-1/1/0.0
```

   If a currently used tunnel interface is deleted from the pool of interfaces, the active mirroring sessions are redistributed from the deleted interface to other tunnel interfaces in the pool. Also, when a new tunnel interface is added into the pool, the service adds the new interface to the list of interfaces available for new mirroring sessions or for existing sessions transferred from a failed interface.

3. Specify the source IP address that the radius-flow-tap service uses for mirroring. This address is used in the IP header prepended to mirrored packets that are sent to the content destination device.

```
[edit services radius-flow-tap]
user@host# set source-ipv4-address ipv4-address
```

4. (Optional) Specify the forwarding class that is applied to the mirrored packets sent to the mediation device.

 Copyright © 2013, Juniper Networks, Inc.

If you do not specify a forwarding class, mirrored packets inherit the forwarding class from the original packet (which is the forwarding class set by default classification that CoS applies to the packet on the ingress interface).

```
[edit services radius-flow-tap]
user@host# set forwarding-class class-name
```

5. (Optional) Specify the lawful intercept policy that determines what traffic, if any, is not sent to the mediation device.

   You can add or change a lawful intercept policy any time, but a changed policy does not apply to a currently enabled policy. To change a policy, add a policy with a new name, use DTCP DISABLE to turn off the current policy, and use DTCP ENABLE to point to the new policy name.

```
[edit services radius-flow-tap]
user@host# set policy policy-name
```

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281
- Guidelines for Configuring Subscriber Secure Policy Mirroring on page 1262

Configuring RADIUS Server Support for Subscriber Secure Policy Mirroring

This topic describes how to configure support for the RADIUS server that initiates subscriber-based traffic mirroring. You create an access profile to specify the RADIUS server support.

To configure the router's interaction with the RADIUS server in support of subscriber secure policy mirroring:

1. Create the access profile and assign a name.

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

2. Specify RADIUS as the authentication method.

```
[edit access profile profile-name]
user@host# set authentication-order radius
```

3. Specify the IP address of the RADIUS server that performs authentication. This server also performs dynamic request (CoA) functions.

```
[edit access profile profile-name]
user@host# set radius authentication-server ip-address
```

4. Specify the secret to use when communicating with the RADIUS server.

```
[edit access profile profile-name]
user@host# set radius-server server-address secret password
```

5. Specify other optional RADIUS configuration settings as needed, such as accounting support.

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- RADIUS Attributes Used for Subscriber Secure Policy on page 1254

Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic

This topic describes the steps to enable subscriber secure policy mirroring of IPv4 multicast traffic. You can enable and disable IPv4 multicast intercept on a per chassis basis.

To configure the radius-flow-tap service to support subscriber secure policy mirroring:

1. Configure the flow-tap service used for subscriber secure policy mirroring.

   ```
   [edit services]
   user@host# edit radius-flow-tap
   ```

2. Enable the interception of multicast traffic.

   ```
   [edit services radius-flow-tap]
   user@host# set multicast-interception
   ```

Related Documentation

- Subscriber Secure Policy Support for IPv4 Multicast Traffic on page 1253
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281

Terminating RADIUS-Initiated Subscriber Traffic Mirroring

You can terminate RADIUS-initiated traffic mirroring sessions by the following action:

- RADIUS CoA message receipt—Terminated upon receipt of a CoA message with the VSA 26-58 (LI-Action) value of 0. The RADIUS administrator configures the LI-Action of 0 in the mirrored subscriber's RADIUS record.

Related Documentation

- RADIUS-Initiated Subscriber Secure Policy Overview on page 1247
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 75

Overview of Subscriber Secure Policy Using DTCP

- DTCP-Initiated Subscriber Secure Policy Overview on page 1267
- Subscriber Secure Policy Traffic Mirroring Architecture Using DTCP on page 1267
- DTCP-Initiated Traffic Mirroring Interfaces on page 1269
- DTCP-Initiated Traffic Mirroring Process on page 1271
- Subscriber Secure Policy Support for IPv4 Multicast Traffic on page 1272
- Subscriber Secure Policy and L2TP LNS Subscribers on page 1272
- DTCP Messages Used for Subscriber Secure Policy on page 1273
- DTCP Traffic Mirroring Triggers on page 1273
- Packet Header for Mirrored Traffic Sent to Mediation Device on page 1278

DTCP-Initiated Subscriber Secure Policy Overview

Dynamic Tasking Control Protocol (DTCP)-initiated mirroring creates secure policies to mirror traffic for the subscriber based on DTCP messages. The attributes in a DTCP ADD message trigger the router to start mirroring traffic and specify the interface on which the mirroring takes place. The mirroring operations can be initiated by DTCP messages as follows:

- Subscriber login—Mirroring starts on the specified interface when the subscriber logs in. The DTCP ADD message must be sent to the router before the subscriber logs in.
- In-session—Mirroring starts for all subscribers that match the trigger supplied in the DTCP ADD message when the router receives a DTCP ADD message.

Related Documentation

- Subscriber Secure Policy Traffic Mirroring Architecture Using DTCP on page 1267
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281

Subscriber Secure Policy Traffic Mirroring Architecture Using DTCP

Figure 45 on page 1268 shows the architecture of the DTCP-initiated subscriber secure policy mirroring environment.

Copyright © 2013, Juniper Networks, Inc.

Figure 45: DTCP-Initiated Subscriber Secure Policy Architecture

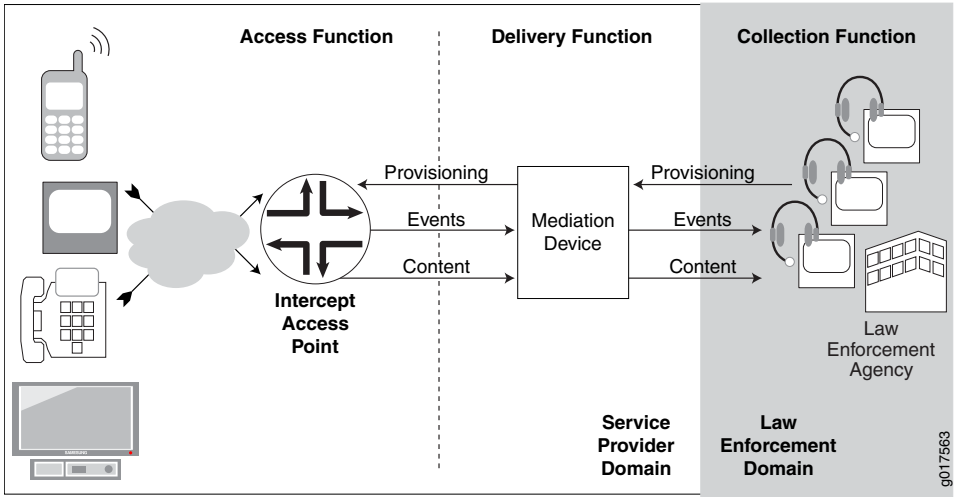

Table 114 on page 1268 describes the functions and components of a DTCP-initiated subscriber secure policy traffic mirroring environment.

Table 114: DTCP-Initiated Subscriber Secure Policy Functions and Components

Function or Component	Description
Collection function	The collection function is responsible for collecting intercepted content and identifying information from the delivery function. The collection function is the responsibility of the law-enforcement agency (LEA).
Delivery function	The delivery function delivers information that it receives from the access function to the collection function. The delivery function is performed by the mediation device.
Access function	The access function has access to the intercept target's traffic content and intercept-related events. It is responsible for collecting this information and sending it to the delivery function. The access function is performed by intercept access points (IAPs).
Events	Intercept-related events, such as login or logout events or mirroring session activation or deactivation. The router sends the events to the mediation device in SNMP traps.
LEA	Law enforcement agency. The LEA provides intercept targets to the service provider who provisions the mediation device.

Copyright © 2013, Juniper Networks, Inc.

Table 114: DTCP-Initiated Subscriber Secure Policy Functions and Components *(continued)*

Function or Component	Description
Mediation device	The mediation device receives provisioning information from the LEA, and it uses the information to send provisioning information to the IAP (the router). The mediation device also receives intercept-related events and intercepted content from the router, and delivers the events and content to the LEA.
IAP	Intercept access point. In a subscriber access network the Juniper Networks router is the IAP. Using subscriber secure policies, the IAP intercepts traffic to and from the subscriber whose traffic is being mirrored. It encapsulates the intercepted content in a packet header and delivers it to the mediation device, while also sending the traffic to the intended destination. The IAP also sends intercept-related events to the mediation device using SNMP traps.

Related Documentation

- DTCP-Initiated Subscriber Secure Policy Overview on page 1267
- DTCP-Initiated Traffic Mirroring Interfaces on page 1269
- DTCP-Initiated Traffic Mirroring Process on page 1271

DTCP-Initiated Traffic Mirroring Interfaces

Figure 46 on page 1269 shows the interfaces involved in DTCP-initiated secure subscriber policy traffic mirroring.

Figure 46: DTCP-Initiated Traffic Mirroring Interfaces

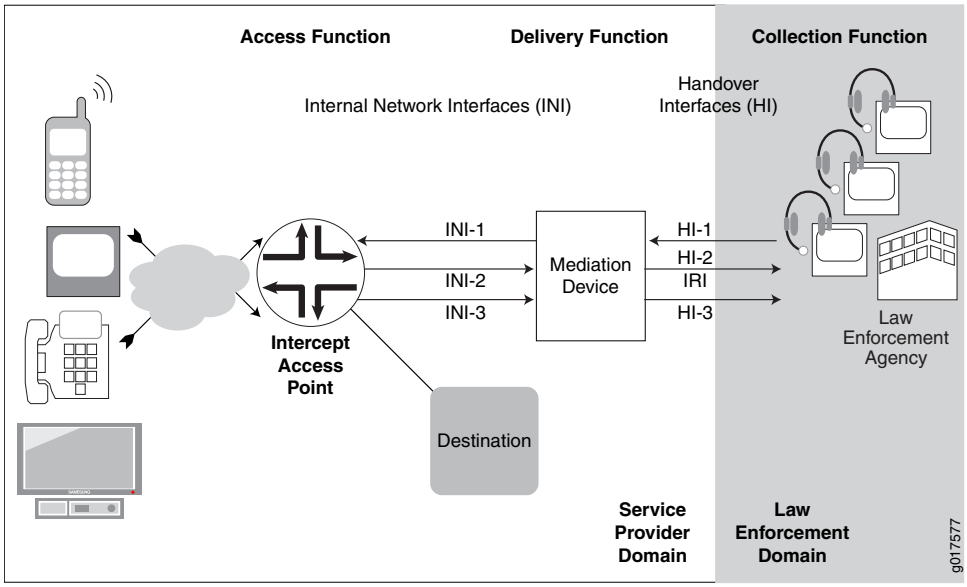

Table 115 on page 1270 describes the interfaces involved in DTCP-initiated secure subscriber policy traffic mirroring.

Table 115: DTCP-Initiated Traffic Mirroring Interfaces

Interface	Description
HI-1	Handover Interface 1—Administrative interface between the LEA and the service provider mediation device. The LEA sends provisioning information to the mediation device on this interface.
HI-2	Handover Interface 2—Intercept-related information interface between the LEA and the mediation device that is used to deliver intercept-related events to the LEA. These events can be subscriber session events such as login, logout, and authentication.
HI-3	Handover Interface 3—Intercepted content Interface between the mediation device and LEA that is used to deliver intercepted content to the LEA.
INI-1	Internal network Interface 1—Interface used to send DTCP messages containing intercept provisioning information from the mediation device to the router.
INI-2	Internal network interface 2—Interface used to send intercept-related events from the router to the mediation device. This information is sent in SNMP traps.
INI-3	Internal network interface 3—Interface used to send intercepted content from the router to the mediation device.

Related Documentation

- Subscriber Secure Policy Traffic Mirroring Architecture Using DTCP on page 1267
- DTCP-Initiated Traffic Mirroring Process on page 1271

 Copyright © 2013, Juniper Networks, Inc.

DTCP-Initiated Traffic Mirroring Process

Figure 47 on page 1271 shows the process for a DTCP-initiated subscriber mirroring operation.

Figure 47: DTCP-Initiated Subscriber Secure Policy Model

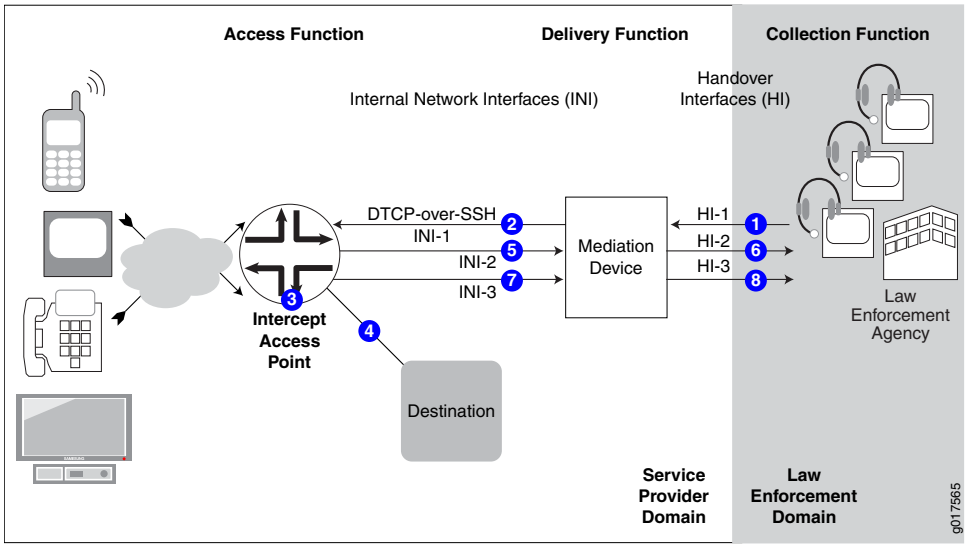

1—The LEA sends provisioning information for a subscriber whose traffic is to be mirrored over the HI-1 interface to the mediation device.	5—As intercept-related events occur, the IAP sends the events in SNMP traps over the INI-2 interface to the mediation device.
2—The mediation device sends a DTCP ADD message that contains provisioning information over the INI-1 interface to the IAP (the router).	6—The mediation device provides the intercept-related events over the HI-2 interface to the LEA.
3—The IAP creates a subscriber secure policy based on information in the DTCP ADD message. If the IAP receives the DTCP ADD before the subscriber logs in, mirroring begins when the subscriber logs in. If the router receives the DTCP ADD after the subscriber logs in, mirroring begins when the ADD message is received.	7—The IAP sends the mirrored content to the mediation device over the INI-3 interface.
4—The IAP sends the original subscriber traffic to its intended destination.	8—The mediation device sends mirrored content over the HI-3 interface to the LEA.

Related Documentation

- Subscriber Secure Policy Traffic Mirroring Architecture Using DTCP on page 1267
- DTCP-Initiated Traffic Mirroring Interfaces on page 1269
- DTCP Messages Used for Subscriber Secure Policy on page 1273
- DTCP Traffic Mirroring Triggers on page 1273

Copyright © 2013, Juniper Networks, Inc.

Subscriber Secure Policy Support for IPv4 Multicast Traffic

IP multicast traffic is used for applications such as audio or video streaming, IPTV, video conferencing, or online gaming. Multicast traffic is sent to multiple subscribers who have joined a multicast group.

Secure subscriber policy allows for the mirroring of IPv4 multicast traffic sent to a specific subscriber. If multiple subscribers whose traffic requires mirroring join the same multicast session, the subscriber secure policy feature mirrors each subscriber's traffic and forwards it separately to the mediation device with the proper prepended header.

Mirroring of multicast traffic is supported only for subscribers in the default logical system.

You can enable and disable the mirroring of multicast traffic on a per-chassis basis. You cannot enable or disable it on a per-subscriber basis.

Triggering the Mirroring of IPv4 Multicast Traffic

Multicast traffic being sent towards a subscriber does not contain much of the identifying information used to trigger mirroring of a subscriber's unicast traffic. For example, the multicast packet contains the multicast group address in the destination address of the packet instead of the subscriber's IP address. It also does not contain the user name or MAC address of the subscriber, and does not include information obtained by RADIUS or DHCP. Therefore, methods of identifying multicast traffic that is received by a subscriber are not the same as methods of identifying a subscriber's unicast traffic or multicast traffic that is sent by a subscriber.

To join a multicast group, a subscriber sends an IGMP join request, and it receives a reply. The reply contains the multicast groups to which the subscriber is registered. Triggering the mirroring of multicast traffic is based on the sending of the IGMP join request and the information in the IGMP reply. If the subscriber's unicast traffic is already being mirrored either through DTCP-initiated or RADIUS-initiated traffic mirroring, and the subscriber sends an IGMP join request, mirroring of multicast traffic sent to the subscriber is initiated. The traffic being mirrored is based on the groups contained in the IGMP reply.

Related Documentation

- Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic on page 1266

Subscriber Secure Policy and L2TP LNS Subscribers

Dynamic Tasking Control Protocol (DTCP)-initiated and RADIUS-initiated per-subscriber traffic mirroring can be applied to Point-to-Point Protocol (PPP) subscribers whose traffic is tunneled with Layer 2 Tunneling Protocol (L2TP). At the L2TP network server (LNS), both subscriber ingress traffic (from the L2TP access concentrator, or LAC, to the LNS) and subscriber egress traffic (from the LNS to the LAC) are mirrored at the inline services (si) interface corresponding to the subscriber. Ingress traffic is mirrored after decapsulation of L2TP, HDLC, and PPP headers. The egress traffic is mirrored before the IP datagram is encapsulated. The mirrored traffic contains only the IP datagram belonging to the subscriber.

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- RADIUS Attributes Used for Subscriber Secure Policy on page 1254

DTCP Messages Used for Subscriber Secure Policy

You can use DTCP to provision traffic mirroring on the router by sending DTCP messages from the mediation device to the router.

There are three types of DTCP messages:

- ADD—Triggers mirroring of subscriber secure policy sessions. You include an attribute that triggers the router to begin mirroring a subscriber session. You can also include attributes that identify where to send the mirrored session data and how to uniquely identify traffic when simultaneous intercepts are active. The ADD message also provides instructions to populate fields in the encapsulation header for packets sent to the mediation device.
- LIST—Requests information about sessions that are currently being mirrored. This information is returned in a LIST response.
- DELETE—Removes a subscriber mirroring trigger or can be used to disable all mirroring.

Related Documentation

- DTCP-Initiated Traffic Mirroring Process on page 1271
- DTCP Traffic Mirroring Triggers on page 1273
- **ADD DTCP on page 1292**
- **DELETE DTCP on page 1295**
- **LIST DTCP on page 1299**

DTCP Traffic Mirroring Triggers

Table 116 on page 1274 lists the DTCP attributes that you can use in DTCP ADD messages to trigger traffic mirroring.

Table 116: DTCP Mirroring Triggers for Use in ADD Messages

Attribute Name	DTCP Message Semantic	Description
Account Session ID	X-Act-Sess-Id	Trigger that is based on the text string of the Account Session ID associated with the subscriber session. If the subscriber logs out, the intercept terminates. We recommend that you use other triggers to ensure that all sessions for a subscriber are intercepted.
Calling Station ID	X-Call-Sta-Id	Trigger that is based on the text string of the Calling Station ID associated with the subscriber. If the subscriber is not logged on, the policy is applied at any current or subsequent subscriber log in.
Drop Policy Name	X-Drop-Policy	Trigger that is based on the name of the configured lawful intercept policy.
IP Address	X-IP-Addr	Trigger for the IPv4 address that is associated with a subscriber. If you use the IP Address trigger, and the subscriber is not using the default logical system, you must include the Logical System attribute in your DTCP message. If the subscriber is not using the default routing instance, you must include the Routing Instance attribute in your DTCP message.
Interface Identifier	X-Interface-Id	Trigger for subscribers that are configured to use a specific router interface. All subscribers that use the interface have their traffic mirrored. Add this attribute as a text string that identifies the physical interface; for example, **ge-0/0/0.1** or **demux0.107472834**.

Copyright © 2013, Juniper Networks, Inc.

Table 116: DTCP Mirroring Triggers for Use in ADD Messages *(continued)*

Attribute Name	DTCP Message Semantic	Description
NAS Port ID	X-NAS-Port-Id	Trigger that is based on the NAS port ID of the subscriber.
Remote Circuit ID	X-RM-Circuit-Id	For DHCP subscribers, trigger that is used with the Remote Agent ID to specify the DHCP option 82 that is associated with this session to completely specify a trigger. For PPPoE subscribers, agent circuit ID (ACI) in the PPPoE Intermediate Agent (PPPoE IA) tag.
Remote Agent ID	X-RM-Agent-Id	For DHCP subscribers, trigger that is used with the Remote Circuit ID to specify the session or by itself to completely specify the trigger. For PPPoE subscribers, agent remote identifier (ARI) in the PPPoE Intermediate Agent (PPPoE IA) tag.
Logical System	X-Logical-System	Trigger attribute that you can use with the IP Address or Subscriber User Name triggers. It is ignored for other triggers. The value **default** is used if no logical system exists for the subscriber.
Routing Instance	X-Router-Instance	Trigger attribute that you can use with the IP Address or Subscriber User Name triggers. It is ignored for other triggers. The value **default** is used if no routing instance exists for the subscriber.

Copyright © 2013, Juniper Networks, Inc.

Table 116: DTCP Mirroring Triggers for Use in ADD Messages *(continued)*

Attribute Name	DTCP Message Semantic	Description
Subscriber User Name	X-UserName	Trigger based on a subscriber username. If you use the Subscriber User Name trigger, and the subscriber is not using the default logical system, you must include the Logical System attribute in your DTCP message. If the subscriber is not using the default routing instance, you must include the Routing Instance attribute in your DTCP message.

Triggering Subscriber Secure Policy for Subscribers on Dynamic Authenticated VLANs

BEST PRACTICE: When you have DHCPv4/DHCPv6 subscribers over VLANs, two sessions are created for each subscriber—one for the Layer 2 VLAN, and one for DHCP. In this case do not use a trigger, such as Remote Circuit ID (ACI), that applies to both the VLAN and the DHCP sessions. If the DHCP and VLAN sessions match the same trigger, the DHCP subscriber login fails and subscriber secure policy is not triggered. You need to select a traffic mirroring trigger that matches only one of these sessions.

Order in Which Trigger Attributes Are Processed

If a subscriber matches more than one of the DTCP mirroring triggers, the router processes mirroring triggers in ADD messages in the following order:

1. Account Session ID
2. Calling Station ID
3. IP Address
4. Interface Identifier
5. NAS Port ID
6. Remote Agent ID
7. Subscriber User Name
8. Drop Policy Name

Related Documentation

- Packet Header for Mirrored Traffic Sent to Mediation Device on page 1278
- ADD DTCP on page 1292
- DELETE DTCP on page 1295

Copyright © 2013, Juniper Networks, Inc.

- LIST DTCP on page 1299
- Example: Using DTCP Messages to Trigger, Verify, and Disable Traffic Mirroring for Subscribers on page 1300

Packet Header for Mirrored Traffic Sent to Mediation Device

When the router sends mirrored traffic to the mediation device, it encapsulates the mirrored payload in a packet header before it sends the mirrored traffic to the mediation device.

Figure 48 on page 1278 is the mirrored packet header that the router sends to the mediation device.

Figure 48: Mirrored Packet Header

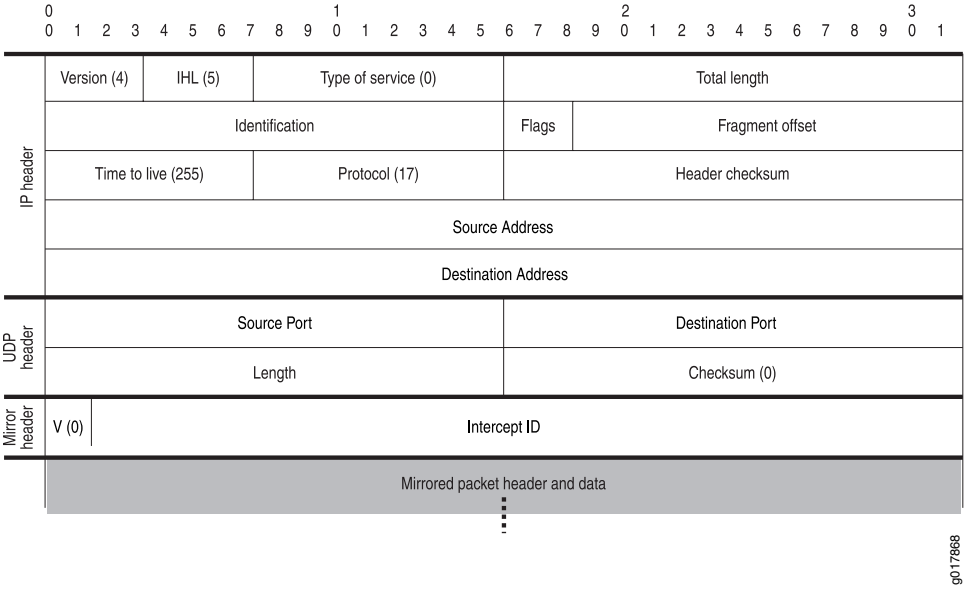

Table 117 on page 1278 describes the fields in the packet header of mirrored packets.

Table 117: Mirrored Packet Header and Payload Field Descriptions

Field	Value	Length (Bits)
IP Header		
Version	4	4
IHL	5	4
Type of Service	0	8
Total Length	Dynamically computed	16
Identification	Dynamically computed	16
Flags	Dynamically computed	3
Fragment Offset	Dynamically computed	13

Copyright © 2013, Juniper Networks, Inc.

Table 117: Mirrored Packet Header and Payload Field Descriptions *(continued)*

Field	Value	Length (Bits)
Time to Live	255	8
Protocol	17	8
Header Checksum	Dynamically computed	16
Source Address	IP address of the router interface that sends mirrored traffic to the mediation device	32
Destination Address	IP address of the mediation device to which mirrored traffic is forwarded. This value is taken from the X-JTap-Cdest-Dest-Address attribute that is sent to the router in the DTCP ADD command.	32
UDP Header		
Source Port	UDP port number on the router from which mirrored traffic is sent to the mediation device	16
Destination Port	UDP port on the mediation device to which mirrored traffic is forwarded. This value is taken from the X-JTap-Cdest-Dest-Port attribute that is sent to the router in the DTCP ADD command.	16
Length	Dynamically computed	16
Checksum	0	16
Mirror Header		
V (mirror header value)	0	2
Intercept ID	Value of the X-MD-Intercept-Id that is sent to the router in the DTCP ADD command.	30

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- DTCP-Initiated Subscriber Secure Policy Overview on page 1267
- **ADD DTCP on page 1292**
- Example: Using DTCP Messages to Trigger, Verify, and Disable Traffic Mirroring for Subscribers on page 1300

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 76

Configuring DTCP-Initiated Subscriber Secure Policy Mirroring

- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281
- Guidelines for Configuring Subscriber Secure Policy Mirroring on page 1282
- Configuring Tunnel Interfaces for Subscriber Secure Policy Mirroring on page 1283
- Configuring Support for Subscriber Secure Policy Mirroring on page 1284
- Example: Configuring Traffic That Is Mirrored Using DTCP-Initiated Subscriber Secure Policy on page 1285
- Configuring the Mediation Device as a User on the Router on page 1287
- Configuring a DTCP-over-SSH Connection to the Mediation Device on page 1288
- Configuring the Mediation Device to Provision Traffic Mirroring on page 1288
- Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic on page 1289
- Terminating DTCP-Initiated Subscriber Traffic Mirroring Sessions on page 1289

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 tunnel interfaces that are used to send mirrored content to the mediation device.

   See "Configuring Tunnel Interfaces for Subscriber Secure Policy Mirroring" on page 1263.

2. Configure the radius-flow-tap service support for secure subscriber policy. This support includes configuring the tunnels and optional forwarding-class information that the

Copyright © 2013, Juniper Networks, Inc.

subscriber secure policy service uses to send mirrored traffic to the content destination device.

See "Configuring Support for Subscriber Secure Policy Mirroring" on page 1264.

3. 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" on page 1287.

4. Configure the mediation device to provision traffic mirroring on the router.

   See "Configuring the Mediation Device to Provision Traffic Mirroring" on page 1288.

5. Configure a DTCP-over-SSH connection to the mediation device.

   See "Configuring a DTCP-over-SSH Connection to the Mediation Device" on page 1288.

6. (Optional) Enable mirroring of IPv4 multicast traffic on the router.

   See "Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic" on page 1266

7. Configure SNMPv3 trap support to report mirroring information to an external device.

   See "Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring" on page 1307.

You can terminate an active subscriber mirroring session at any time.

See "Terminating DTCP-Initiated Subscriber Traffic Mirroring Sessions" on page 1289.

Related Documentation

- DTCP-Initiated Subscriber Secure Policy Overview on page 1267
- Intercept-Related Events Transmitted to the Mediation Device on page 1305

Guidelines for Configuring Subscriber Secure Policy Mirroring

The subscriber secure policy service uses the radius-flow-tap service infrastructure.

When configuring subscriber secure policy mirroring, consider the following guidelines regarding the relationship between subscriber secure policy service and the radius-flow-tap service:

- The radius-flow-tap service **[edit services radius-flow-tap]** and the flow-tap service **[edit services flow-tap]** cannot run simultaneously on the router. Therefore, flow-tap and subscriber secure policy mirroring cannot run simultaneously on the same router.
- You can configure one instance of the radius-flow-tap service on the router. Subscriber secure policy RADIUS-initiated mirroring and DTCP-initiated mirroring use the radius-flow-tap service.
- If you delete the radius-flow-tap service all subscriber secure policy mirroring stops.

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281

Copyright © 2013, Juniper Networks, Inc.

• Configuring Support for Subscriber Secure Policy Mirroring on page 1264

Configuring Tunnel Interfaces for Subscriber Secure Policy Mirroring

The router, acting as the IAP, uses tunnel interfaces (vt interfaces) to send mirrored traffic to the mediation device. The IAP equally distributes the mirrored traffic across the available tunnel interfaces.

Because the MX Series 3D Universal Edge Routers do not support Tunnel Services PICs, you create a pool tunnel interfaces on MX Series routers at the **[edit chassis]** hierarchy level.

You can configure up to 2048 mirrored subscriber sessions per chassis.

To configure a pool of tunnel interfaces for use by subscriber secure policy mirroring:

1. Access the chassis configuration, and specify the slot number of the DPC, MPC, or MIC.
   - On the MX80 router, the range is 0 through 1.
   - On other MX Series routers, if two System Control Boards (SCBs), are installed, the range is 0 through 11. If three SCBs are installed, the range is 0 through 5 and 7 through 11.

     ```
     [edit chassis]
     user@host# edit fpc 1
     ```

2. Configure the PIC number of the FPC.
   - On MX80 routers, if the FPC is 0, the PIC number can only be 0. If the FPC is 1, the PIC range is 0 through 3.
   - For all other MX Series routers, the range is 0 through 3.

     ```
     [edit chassis fpc 1]
     user@host# edit pic 1
     ```

3. Specify that the FPC and PIC are to be used for tunnel interfaces.

   ```
   [edit chassis fpc 1 pic 1]
   user@host# edit tunnel-services
   ```

4. Specify the amount of bandwidth to reserve for tunnel traffic on each Packet Forwarding Engine.
   - 1g indicates that 1 Gbps of bandwidth is reserved for tunnel traffic.
   - 10g indicates that 10 Gbps of bandwidth is reserved for tunnel traffic.

   If you specify a bandwidth that is not compatible, tunnel services are not activated. For example, you cannot specify a bandwidth of 1 Gbps for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC.

   ```
   [edit chassis fpc 1 pic 1 tunnel-services]
   user@host#
   ```

```
user@host# set bandwidth 1g
```

5. Configure the tunnel interfaces, including the family.

   To configure subscriber secure policy mirroring for IPv6 traffic, configure the tunnel interfaces for both the **inet** and **inet6** families.

```
[edit interfaces]
user@host# set vt-1/1/0 unit 0 family inet
user@host# set vt-1/1/0 unit 0 family inet6
```

Related Documentation

- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281
- Subscriber Secure Policy and L2TP LNS Subscribers on page 1259

Configuring Support for Subscriber Secure Policy Mirroring

Subscriber secure policy runs on the radius-flow-tap service. This topic describes the steps to configure radius-flow-tap support for RADIUS-initiated and DTCP-initiated subscriber secure policy mirroring.

To configure the radius-flow-tap service to support subscriber secure policy mirroring:

1. Configure the flow-tap service used for subscriber secure policy mirroring.

```
[edit services]
user@host# edit radius-flow-tap
```

2. Assign the tunnel interfaces that the radius-flow-tap service uses.

```
[edit services radius-flow-tap]
user@host# set interfaces vt-1/1/0.0
```

   If a currently used tunnel interface is deleted from the pool of interfaces, the active mirroring sessions are redistributed from the deleted interface to other tunnel interfaces in the pool. Also, when a new tunnel interface is added into the pool, the service adds the new interface to the list of interfaces available for new mirroring sessions or for existing sessions transferred from a failed interface.

3. Specify the source IP address that the radius-flow-tap service uses for mirroring. This address is used in the IP header prepended to mirrored packets that are sent to the content destination device.

```
[edit services radius-flow-tap]
user@host# set source-ipv4-address ipv4-address
```

4. (Optional) Specify the forwarding class that is applied to the mirrored packets sent to the mediation device.

   If you do not specify a forwarding class, mirrored packets inherit the forwarding class from the original packet (which is the forwarding class set by default classification that CoS applies to the packet on the ingress interface).

```
[edit services radius-flow-tap]
user@host# set forwarding-class class-name
```

 Copyright © 2013, Juniper Networks, Inc.

5. (Optional) Specify the lawful intercept policy that determines what traffic, if any, is not sent to the mediation device.

   You can add or change a lawful intercept policy any time, but a changed policy does not apply to a currently enabled policy. To change a policy, add a policy with a new name, use DTCP DISABLE to turn off the current policy, and use DTCP ENABLE to point to the new policy name.

   ```
   [edit services radius-flow-tap]
   user@host# set policy policy-name
   ```

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281
- Guidelines for Configuring Subscriber Secure Policy Mirroring on page 1262

Example: Configuring Traffic That Is Mirrored Using DTCP-Initiated Subscriber Secure Policy

This example shows how to configure traffic that is mirrored using DTCP-initiated subscriber secure policy.

- Requirements on page 1285
- Overview on page 1285
- Configuration on page 1286

Requirements

- Juniper Networks MX Series routers.
- Junos OS Release 12.3R1 or later.

Overview

This example drops all video on demand TCP traffic from subnet 10.0.0.0/8 to any subscriber on which the policy named vod is enabled.

To configure traffic mirroring using DTCP-initiated subscriber secure policy:

1. Create a policy.
2. Set up the policy to filter IPv4 or IPv6 traffic by source or destination address, or port, protocol, or DSCP value.
3. Apply the policy using the DTCP attribute X-Drop-Policy.
4. Use the X-Drop-Policy with the ADD DTCP command to begin filtering traffic when mirroring is triggered.

NOTE: To begin filtering traffic that is currently being mirrored, use the X-Drop-Policy attribute with the new ENABLE DTCP command. To stop filtering traffic that is currently being mirrored, use the X-Drop-Policy attribute with the new DISABLE DTCP command.

Configuration

Step-by-Step Procedure

To configure filtering mirrored traffic before it is sent to a mediation device:

1. Specify that you want to configure radius-flow-tap.

   ```
   [edit services]
   user@host# edit radius-flow-tap
   ```

2. Specify that you want to configure a video on demand policy.

   ```
   [edit services radius-flow-tap]
   user@host# edit policy vod
   ```

3. Specify inet as the family that you want to use.

   ```
   [edit services radius-flow-tap vod]
   user@host# edit inet
   ```

4. Specify t1 as the term name for the IPv4 drop-policy.

   ```
   [edit services radius-flow-tap vod inet]
   user@host# edit drop-policy t1
   ```

5. Specify the source address for the drop-policy.

   ```
   [edit services radius-flow-tap vod inet drop-policy t1]
   user@host# edit source-address 10.0.0.0/8
   ```

6. Specify the match criteria that you want to use.

   ```
   [edit services radius-flow-tap vod inet drop-policy t1]
   user@host# set protocol tcp
   ```

Results

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

```
[edit services radius-flow-tap policy]
vod {
  inet {
    drop-policy t1 {
      from{
        source-address {
          10.0.0.0/8;
        }
        protocol tcp;
      }
    }
  }
```

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

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring Support for Subscriber Secure Policy Mirroring on page 1264
- DTCP Traffic Mirroring Triggers on page 1273

Configuring the Mediation Device as a User on the Router

In order for the router to receive DTCP messages from the mediation device, you need to configure the mediation device as a user on the router. To do so, create a login class that provides flow-tap operation permission and then create a login account that uses the login class.

To configure the mediation device as a user on the router:

1. Create the login class and configure **flow-tap-operation** permissions for the class.
   a. Specify that you want to configure login properties.

      [edit system]
      user@host# **edit login**

   b. Create and name the class.

      [edit system login]
      user@host# **edit class** *class-name*

   c. Configure the **flow-tap-operation** permission for the class.

      [edit system login class *class-name*]
      user@host# **set permissions flow-tap-operation**

2. Create the user login account for the mediation device.
   a. Create the user account.

      [edit system login]
      user@host# **edit user** *username*

   b. Configure the user ID.

      [edit system login user *username*]
      user@host# **set uid** *uid-value*

   c. Configure the class for the user account.

      [edit system login user *username*]
      user@host# **set class** *class-name*

   d. Configure the authentication for the user account.

      [edit system login user *username*]
      user@host# **set authentication encrypted-password** *password*

Configuring a DTCP-over-SSH Connection to the Mediation Device

DTCP-initiated subscriber secure policy requires a DTCP-over-SSH connection for the flow-tap service. This connection is used to send provisioning information from the mediation device to the router.

To enable the DTCP-over-SSH flow-tap service to support subscriber secure policy mirroring:

1. Access the **flow-tap-dtcp** service.

   [edit system services]
   user@host# **edit flow-tap-dtcp**

2. Enable SSH support for DTCP.

   [edit system services flow-tap-dtcp]
   user@host# **set ssh**

3. (Optional) Configure maximum number of established connections allowed for the DTCP service.

   [edit system services flow-tap-service ssh]
   user@host# **set connection-limit *limit***

4. (Optional) Configure the maximum number of connection attempts allowed per minute for DTCP.

   [edit system services flow-tap-service ssh]
   user@host# **set rate-limit *limit***

Related Documentation

- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281

Configuring the Mediation Device to Provision Traffic Mirroring

To set up the mediation device to provision traffic mirroring on the router, use the following DTCP messages:

- To configure traffic-mirroring triggers, use the **ADD DTCP** message.
- To remove an existing traffic-mirroring trigger, use the **DELETE DTCP** message.
- To show existing traffic-mirroring triggers, use the **LIST DTCP** message.

For an example of how to use the DTCP messages, see "Example: Using DTCP Messages to Trigger, Verify, and Disable Traffic Mirroring for Subscribers" on page 1300.

Related Documentation

- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281

Copyright © 2013, Juniper Networks, Inc.

Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic

This topic describes the steps to enable subscriber secure policy mirroring of IPv4 multicast traffic. You can enable and disable IPv4 multicast intercept on a per chassis basis.

To configure the radius-flow-tap service to support subscriber secure policy mirroring:

1. Configure the flow-tap service used for subscriber secure policy mirroring.

   ```
   [edit services]
   user@host# edit radius-flow-tap
   ```

2. Enable the interception of multicast traffic.

   ```
   [edit services radius-flow-tap]
   user@host# set multicast-interception
   ```

Related Documentation

- Subscriber Secure Policy Support for IPv4 Multicast Traffic on page 1253
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281

Terminating DTCP-Initiated Subscriber Traffic Mirroring Sessions

You can terminate DTCP-initiated traffic mirroring sessions by the following action:

- DTCP DELETE message receipt—Terminated upon receipt of a DTCP DELETE message. The DTCP administrator configures the DELETE message to include the same mirroring attributes that are used in the ADD message to initiate mirroring.

Related Documentation

- DELETE DTCP on page 1295
- DTCP Messages Used for Subscriber Secure Policy on page 1273

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 77

DTCP Messages Used for DTCP-Initiated Subscriber Secure Policy Mirroring

- Example: Using DTCP Messages to Trigger, Verify, and Disable Traffic Mirroring for Subscribers on page 1300

ADD DTCP

Syntax

ADD DTCP/0.7
Csource-ID: *user-name*
Cdest-ID: *variable*
Priority: *priority-number*
X-Drop-Policy: *policy-name*
X-JTap-Cdest-Dest-Address: *ipv4-address*
X-JTap-Cdest-Dest-Port: *udp-port*
X-JTap-Cdest-Source-Address: *ipv4-address*
X-JTap-Cdest-Source-Port: *port-number*
X-JTap-Cdest-TTL: *time-to-live*
X-MD-Intercept-Id: *8-byte-id*
Dtcp-trigger: trigger-value
Dtcp-attribute: attribute-value
Flags: *flag*
Seq: *sequence-number*
Authentication-Info: *ssh-authentication-string*

Description

Specify the DTCP attributes used in ADD messages to cause the router to trigger traffic mirroring and provide instructions to populate fields in the encapsulation header for packets sent to the mediation device.

The DTCP ADD message can be sent either before or after subscribers log in through the interface.

The following attributes are added to the packet header of mirrored packets that the router sends to the mediation device. These attributes are required in the DTCP ADD message.

- X-JTap-Cdest-Dest-Address
- X-JTap-Cdest-Dest-Port
- X-MD-Intercept-Id

Options

Csource-ID: *user-name*—Username on the router. This username must be configured as a DTCP user on the router using the **set system login class** or **set system login user** statements.

Cdest-ID: *variable*—ID of the mediation device.

Flags: *flag*—STATIC is the only flag supported.

Priority: *priority-number*—This implementation of DTCP does not use the priority number.

X-Drop-Policy *policy-name*—Name of the policy used to determine which mirrored packets are no longer sent to the mediation device.

X-JTap-Cdest-Dest-Address: *ipv4-address*—Destination IPv4 address of the mediation device to which intercepted packets are sent. You must include this attribute in your ADD messages.. It is used in the header of mirrored traffic that is sent to the mediation device.

Copyright © 2013, Juniper Networks, Inc.

X-JTap-Cdest-Dest-Port: ***udp-port***—Destination port of the mediation device to which intercepted packets are sent. You must include this attribute in your ADD messages. It is used in the header of mirrored traffic that is sent to the mediation device.

X-JTap-Cdest-Source-Address: ***ipv4-address***—Source IPv4 address. You must include this attribute in your ADD messages. If the value entered does not match the value configured on the router using the **set services radius-flow-tap source-ipv4-address** ***source-ipv4-address*** statement, it is replaced by configured value.

X-JTap-Cdest-Source-Port: ***port-number***—Source port. You must include this attribute in your ADD messages. If the value entered does not match the value of X-Jtap-Cdest-Dest-Port, it is ignored.

X-JTap-Cdest-TTL: ***time-to-live***—TTL value to be used in the forwarded packet.

X-MD-Intercept-Id ***8-byte-id***—An Id that is used to identify a subscriber. You must include this attribute in your ADD messages. This ID is used in the header of mirrored traffic that is sent to the mediation device to allow the device to track a subscriber. The X-MD-Intercept-ID attribute must consist of 8-bytes, and the first two bits must be 00.

Dtcp-trigger*: *trigger-value—DTCP attribute used to trigger traffic mirroring. "DTCP Traffic Mirroring Triggers" on page 1273 lists the DTCP attributes that you can use in DTCP ADD messages to trigger traffic mirroring.

Dtcp-attribute*: *attribute-value—DTCP attribute included in the ADD messages. "DTCP Traffic Mirroring Triggers" on page 1273 lists the DTCP attributes that you can use in ADD messages.

Seq: ***sequence-number***—Number added by the mediation device. DTCP messages contain a monotonically increasing sequence number for each successive message.

Authentication-Info: ***ssh-authentication-string***—String used when you are using SSH to connect to the router.

Required Privilege Level

Not applicable.

Related Documentation

- DTCP Traffic Mirroring Triggers on page 1273
- DTCP-Initiated Subscriber Secure Policy Overview on page 1267
- **DELETE DTCP on page 1295**
- **LIST DTCP on page 1299**

Sample Output

```
ADD DTCP/0.7
Csource-ID: ft-user1
Cdest-ID: cd1
Priority: 2
X-JTap-Cdest-Dest-Address: 10.10.2.50
X-JTap-Cdest-Dest-Port: 7890
X-JTap-Cdest-Source-Address: 10.10.2.9
X-JTap-Cdest-Source-Port: 12321
X-Interface-Id: ge-0/0/2.1
X-MD-Intercept-Id: 55667788
Flags: STATIC
Seq: 1
Authentication-Info: c16d2d9d1679facf0c4a66683af6114d341e4033

DTCP/0.7 200 OK
SEQ: 7
CRITERIA-ID: 2
TIMESTAMP: 2011-02-13 15:56:49.609
```


Copyright © 2013, Juniper Networks, Inc.

DELETE DTCP

Syntax

DELETE DTCP/0.7
Csource-ID: *user-name*
CRITERIA-ID: *criteria-id*
Cdest-ID: *variable*
Flags: *flag*
Seq: *sequence-number*
Authentication-Info:*ssh-authentication-string*

Description

Disable traffic mirroring for a subscriber. Mirroring of the existing subscriber is stopped.

Options

Csource-ID: *user-name*—Username on the router. This name must be configured on the router.

CRITERIA-ID: *criteria-id*—ID that DTCP assigns for the mirrored session when you create a DTCP ADD message. Use this ID in your DELETE messages to disable the intercept for a specific subscriber. To view the ID, use the DTCP LIST message. The CRITERIA-ID and the Cdest-ID are mutually exclusive in DELETE messages.

Cdest-ID: *variable*—ID of the mediation device. Use this ID in your DELETE messages to remove all mirroring sessions associated with a mediation device. The Cdest-ID and the CRITERIA-ID are mutually exclusive in DELETE messages.

Flags: *flag*—STATIC is the only flag supported.

Seq: *sequence-number*—Number added by the mediation device. DTCP messages contain a monotonically increasing sequence number for each successive message.

Authentication-Info: *ssh-authentication-string*—String used when you are using SSH to connect to the router.

Required Privilege Level

Not applicable.

Related Documentation

- DTCP Traffic Mirroring Triggers on page 1273
- DTCP-Initiated Subscriber Secure Policy Overview on page 1267
- **ADD DTCP on page 1292**
- **LIST DTCP on page 1299**

List of Sample Output

DELETE DTCP on page 1296

Sample Output

The following sample shows how to disable mirroring for a specific subscriber by using the CRITERIA-ID.

DELETE DTCP

```
DELETE DTCP/0.7
Csource-ID: dtcp1
CRITERIA-ID: 2
Flags: STATIC
Seq: 10
Authentication-Info: 7e84ae871b12f2da023b038774115bb8d955f17e


DTCP/0.7 200 OK
SEQ: 10
CRITERIA-COUNT: 1
TIMESTAMP: 2011-02-13 16:00:02.802
AUTHENTICATION-INFO: 2834ff32ec07d84753a046cfb552e072cc27d50b
```


Copyright © 2013, Juniper Networks, Inc.

DISABLE DTCP

Syntax

DISABLE DTCP/0.7
Csource-ID: *user-name*
Criteria-ID: *variable*
X-Drop-Policy: *variable*
Flags: *flags*

Release Information

Command introduced in Junos OS Release 12.3.

Description

Specify the DTCP ENABLE message to remove a drop policy that exists because of a prior DTCP ADD or DTCP ENABLE command

The DTCP DISABLE message can only be issued on a Criteria-ID that was returned in a response to a previous DTCP ADD. The policy applies to any new subscribers that match the trigger corresponding to the Criteria-ID. Any existing mirroring remains in place, the policy is not be applied to them.

Options

Csource-ID: *user-name*—Username on the router. This username must be configured as a DTCP user on the router using the **set system login class** or **set system login user** statements.

Criteria-ID: *variable*—Identifies the subscriber on which the policy update occurs.

Flags: *flag*—STATIC is the only flag supported.

X-Drop-Policy: *variable*—Name of the policy that determines which mirrored packets are no longer sent to the mediation device.

Required Privilege Level

Not applicable.

Related Documentation

- ENABLE DTCP on page 1298

Sample Output

```
DISABLE DTCP/0.7
Csource-ID: ft-user1
Criteria-ID: 1
X-Drop: T1
Flags: STATIC
Seq: 1
Authentication-Info: c16d2d9d1679facf0c4a66683af6114d341e4033

DTCP/0.7 200 OK
SEQ: 7
CRITERIA-ID: 2
TIMESTAMP: 2011-02-13 15:56:49.609
```

Copyright © 2013, Juniper Networks, Inc.

ENABLE DTCP

Syntax

ENABLE DTCP/0.7
Csource-ID: *user-name*
Criteria-ID: *variable*
X-Drop-Policy: *variable*
Flags: *flags*

Release Information

Command introduced in Junos OS Release 12.3.

Description

Specify the DTCP attributes used in ENABLE messages to cause the router to trigger a drop policy if one does not already exist from a prior DTCP ADD or DTCP ENABLE command.

The DTCP ENABLE message can only be issued on a Criteria-ID that was returned in a response to a previous DTCP ADD command. The policy applies to any new subscribers who match the trigger corresponding to the Criteria-ID. Any existing mirroring remains in place and the policy is not be applied to them. The DTCP ENABLE command stops only the traffic that is identified by the specified policy from being sent to the mediation device.

Options

Csource-ID: *user-name*—Username on the router. This username must be configured as a DTCP user on the router using the **set system login class** or **set system login user** statements.

Criteria-ID: *variable*—Value returned from a prior DTCP ADD that identifies the trigger on which to disable this drop policy.

Flags: *flag*—STATIC is the only flag supported.

X-Drop-Policy: *variable*—Name of the policy that determines which mirrored packets are no longer sent to the mediation device.

Required Privilege Level

Not applicable.

Related Documentation

- DISABLE DTCP on page 1297

Sample Output

```
ENABLE DTCP/0.7
Csource-ID: ft-user1
Criteria-ID: 1
X-Drop: T1
Flags: STATIC
Seq: 1
Authentication-Info: c16d2d9d1679facf0c4a66683af6114d341e4033

DTCP/0.7 200 OK
SEQ: 7
CRITERIA-ID: 2
TIMESTAMP: 2011-02-13 15:56:49.609
```


Copyright © 2013, Juniper Networks, Inc.

LIST DTCP

Syntax

LIST DTCP/0.7
Csource-ID: *user-name*
Cdest-ID: *variable*
Flags: *flags*
Seq: *sequence-number*
Authentication-Info:*ssh-authentication-string*

Description

Request information that is returned in a LIST response. The response lists triggers only. It does not return sessions that are being mirrored.

Options

Csource-ID: ***user-name***—Username on the router. This name must be configured on the router.

Cdest-ID: ***variable***—ID of the mediation device.

Flags: ***flag***—BOTH is the only flag supported.

Seq: ***sequence-number***—Number added by the mediation device. DTCP messages contain a monotonically increasing sequence number for each successive message.

Authentication-Info: ***ssh-authentication-string***—String used when you are using SSH to connect to the router.

Required Privilege Level

Not applicable.

Related Documentation

- DTCP Traffic Mirroring Triggers on page 1273
- DTCP-Initiated Subscriber Secure Policy Overview on page 1267
- **ADD DTCP on page 1292**
- **DELETE DTCP on page 1295**

List of Sample Output

LIST DTCP on page 1300

Sample Output

LIST DTCP

```
LIST DTCP/0.7
Csource-ID: dtcp1
Cdest-ID: cd1
Seq: 9
Authentication-Info: f6dd64643021debb167ce2fb2d3c7b6622a87e09

DTCP/0.7 200 OK
SEQ: 9
TIMESTAMP: 2011-02-13 15:57:47.667
CRITERIA-ID: 2
CSOURCE-ID: dtcp1
CDEST-ID: cd1
CSOURCE-ADDRESS: 10.10.4.224
FLAGS: Static
X-JTAP-CDEST-DEST-ADDRESS: 192.0.40.168
X-JTAP-CDEST-DEST-PORT: 65535
X-JTAP-CDEST-SOURCE-ADDRESS: 198.15.0.10
X-JTAP-CDEST-SOURCE-PORT: 50000
X-JTAP-CDEST-TTL: 64
X-INTERFACE-ID: demux0.30010002
X-MD-INTERCEPT-ID: 0x0101010130010002
CRITERIA-NUM: 1
CRITERIA-COUNT: 0

CRITERIA-ID: 3
CSOURCE-ID: dtcp1
CDEST-ID: cd1
CSOURCE-ADDRESS: 10.10.4.224
FLAGS: Static
X-JTAP-CDEST-DEST-ADDRESS: 192.0.40.168
X-JTAP-CDEST-DEST-PORT: 65535
X-JTAP-CDEST-SOURCE-ADDRESS: 198.15.0.10
X-JTAP-CDEST-SOURCE-PORT: 50000
X-JTAP-CDEST-TTL: 64
X-INTERFACE-ID: demux0.30010001
X-MD-INTERCEPT-ID: 0x0101010130010001
CRITERIA-NUM: 2
CRITERIA-COUNT: 2
AUTHENTICATION-INFO: 361171ccb24dde6afe8ef66021287f9b8ac16028
```

Example: Using DTCP Messages to Trigger, Verify, and Disable Traffic Mirroring for Subscribers

This example shows how to create DTCP messages to do the following:

- Trigger traffic mirroring for two subscribers based on interface ID.
- Trigger a drop policy if one does not already exist.
- Remove an existing drop policy.
- Verify that subscriber traffic on the two interfaces is being mirrored.
- Disable traffic mirroring on the two subscriber interfaces.
- Verify that traffic mirroring was stopped on the two subscriber interfaces.

Copyright © 2013, Juniper Networks, Inc.

In this example, SSH is being used to communicate with the router.

Creating DTCP ADD Messages to Trigger Traffic Mirroring

This section shows examples of DTCP ADD messages on a mediation device that use the interface ID to trigger traffic mirroring on interfaces demux0.30010002 and demux0.30010001.

```
ADD DTCP/0.7
Csource-ID: dtcp1
Cdest-ID: cd1
Priority: 2
X-JTap-Cdest-Dest-Address: 192.0.40.168
X-JTap-Cdest-Dest-Port: 65535
X-JTap-Cdest-Source-Address: 198.15.0.10
X-JTap-Cdest-Source-Port: 50000
X-JTap-Cdest-TTL: 64
X-Interface-Id: demux0.30010002     /*Used as trigger*/
X-MD-Intercept-Id: 0x0101010130010002
Flags: STATIC
Seq: 7
Authentication-Info: c16d2d9d1679facf0c4a66683af6114d341e4033


DTCP/0.7 200 OK
SEQ: 7
CRITERIA-ID: 2
TIMESTAMP: 2011-02-13 15:56:49.609
AUTHENTICATION-INFO: 4880de4b8cead98c95813fd9b95e240b107d4693


ADD DTCP/0.7
Csource-ID: dtcp1
Cdest-ID: cd1
Priority: 2
X-JTap-Cdest-Dest-Address: 192.0.40.168
X-JTap-Cdest-Dest-Port: 65535
X-JTap-Cdest-Source-Address: 198.15.0.10
X-JTap-Cdest-Source-Port: 50000
X-JTap-Cdest-TTL: 64
X-Interface-Id: demux0.30010001     /*Used as trigger*/
X-MD-Intercept-Id: 0x0101010130010001
Flags: STATIC
Seq: 8
Authentication-Info: dc3c55481a3810c7dd29fdc1b4681d978ff4e7c4

DTCP/0.7 200 OK
SEQ: 8
CRITERIA-ID: 3
TIMESTAMP: 2011-02-13 15:57:20.640
AUTHENTICATION-INFO: 4b31ef1311647e5ba52d2d5d4237b9e5beaa47b7

ADD DTCP/0.7
Csource-ID: ft-user1
Cdest-ID: cd1
Priority: 2
X-JTap-Cdest-Dest-Address:  1.1.1.2
X-JTap-Cdest-Dest-Port: 7899
X-JTap-Cdest-Source-Address: 2.2.2.9
X-JTap-Cdest-Source-Port: 12321
```

```
X-Username: testuser
X-MD-Intercept-Id: 55667789
Flags: STATIC


DTCP/0.7 200 OK
SEQ: 100
CRITERIA-ID: 1
```

Creating DTCP ENABLE Messages to Trigger Traffic Mirroring

This section shows an example of DTCP ENABLE messages on a mediation device that use the interface ID to trigger traffic mirroring on interfaces demux0.30010002 and demux0.30010001.

```
ENABLE DTCP/0.8
Csource-ID: ft-user1
Cdest-ID: cd1
X-Drop-Policy: vod
Flags: STATIC
```

Creating DTCP DISABLE Messages to Trigger Traffic Mirroring

This section shows examples of DTCP DISABLE messages on a mediation device that use the interface ID to trigger traffic mirroring on interfaces demux0.30010002 and demux0.30010001. Whether you used DTCP ADD plus a policy or DTCP ADD and DTCP ENABLE, you can turn the policy off with DTCP DISABLE.

```
DISABLE DTCP/0.8
Csource-ID: ft-user1
Criteria-ID: 1
X-Drop-Policy: vod
Flags: STATIC

DISABLE DTCP/0.8
Csource-ID: ft-user1
Cdest-ID: cd1
X-Drop-Policy: vod
Flags: STATIC
```

Using LIST Messages to Verify That Subscriber Traffic Is Being Mirrored

This section shows examples of a LIST message on the mediation device. The LIST message requests information about the subscribers being mirrored. The information is returned in a LIST response. The response shows that traffic for the two interfaces—demux0.30010002 and demux0.30010001—is being mirrored.

```
LIST DTCP/0.7
Csource-ID: dtcp1
Cdest-ID: cd1
Seq: 9
Authentication-Info: f6dd64643021debb167ce2fb2d3c7b6622a87e09


DTCP/0.7 200 OK
SEQ: 9
TIMESTAMP: 2011-02-13 15:57:47.667
CRITERIA-ID: 2
```

 Copyright © 2013, Juniper Networks, Inc.

```
CSOURCE-ID: dtcp1
CDEST-ID: cd1
CSOURCE-ADDRESS: 10.10.4.224
FLAGS: Static
X-JTAP-CDEST-DEST-ADDRESS: 192.0.40.168
X-JTAP-CDEST-DEST-PORT: 65535
X-JTAP-CDEST-SOURCE-ADDRESS: 198.15.0.10
X-JTAP-CDEST-SOURCE-PORT: 50000
X-JTAP-CDEST-TTL: 64
X-INTERFACE-ID: demux0.30010002    /*subscriber interface*/
X-MD-INTERCEPT-ID: 0x0101010130010002
CRITERIA-NUM: 1
CRITERIA-COUNT: 0

CRITERIA-ID: 3
CSOURCE-ID: dtcp1
CDEST-ID: cd1
CSOURCE-ADDRESS: 10.10.4.224
FLAGS: Static
X-JTAP-CDEST-DEST-ADDRESS: 192.0.40.168
X-JTAP-CDEST-DEST-PORT: 65535
X-JTAP-CDEST-SOURCE-ADDRESS: 198.15.0.10
X-JTAP-CDEST-SOURCE-PORT: 50000
X-JTAP-CDEST-TTL: 64
X-INTERFACE-ID: demux0.30010001   /*subscriber interface*/
X-MD-INTERCEPT-ID: 0x0101010130010001
CRITERIA-NUM: 2
CRITERIA-COUNT: 2
AUTHENTICATION-INFO: 361171ccb24dde6afe8ef66021287f9b8ac16028
```

Using DELETE Messages to Remove Traffic Mirroring Triggers

This section shows examples of DELETE messages used to remove traffic mirroring triggers on demux0.30010001 and demux0.30010002. DTCP DELETE can use either Criteria-ID to delete only that criteria or Cdest-ID to delete everything with cdest-ID that you previously created.

```
DELETE DTCP/0.7
Csource-ID: dtcp1
CRITERIA-ID: 2
Flags: STATIC
Seq: 10
Authentication-Info: 7e84ae871b12f2da023b038774115bb8d955f17e


DTCP/0.7 200 OK
SEQ: 10
CRITERIA-COUNT: 1
TIMESTAMP: 2011-02-13 16:00:02.802
AUTHENTICATION-INFO: 2834ff32ec07d84753a046cfb552e072cc27d50b


DELETE DTCP/0.7
Csource-ID: dtcp1
CRITERIA-ID: 3
Flags: STATIC
Seq: 12
Authentication-Info: 7653fd94659a7183a990bdea654a1b97c0895348

DTCP/0.7 200 OK
```

```
SEQ: 12
CRITERIA-COUNT: 1
TIMESTAMP: 2011-02-13 16:01:35.895
AUTHENTICATION-INFO: 7cd8171057a327434e1b2d9b35f43b88305f9a74
```

Using Disable Messages to Disable Traffic Mirroring Triggers

This section shows an example of

Using Enable Messages to Enable Traffic Mirroring Triggers

This section shows an example of

Verifying That Traffic Mirroring Was Stopped on the Subscriber Interfaces

This section shows an example of a LIST message used to show that traffic mirroring on demux0.30010001 and demux0.30010002 is disabled.

```
LIST DTCP/0.7
Csource-ID: dtcp1
Cdest-ID: cd1
Seq: 13
Authentication-Info: 7c9f825427cfeaecebb0d13ea3842af1021c7d26

DTCP/0.7 430 Unknown Content Destination
SEQ: 13
AUTHENTICATION-INFO: 5ca2eec65106354fe59c878b4c36b7de3c511acd
```

Related Documentation

- DTCP-Initiated Subscriber Secure Policy Overview on page 1267
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 78

Reporting Intercept-Related Events for Subscriber Secure Policy

- Intercept-Related Events Transmitted to the Mediation Device on page 1305
- SNMP Traps for Subscriber Secure Policy LAES Compliance on page 1305
- Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring on page 1307
- Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring on page 1307

Intercept-Related Events Transmitted to the Mediation Device

You can use SNMPv3 traps to report intercept-related events to the mediation device. These events include identifying information for subscribers, such as username or IP address, and subscriber session events, such as login or logout events or mirroring session activation or deactivation. The router sends the events to the mediation device in SNMP traps. Using SNMPv3 provides secure traps that are visible only to authorized individuals on the intended secure mediation device. The traps help support compliance with the Communications Assistance for Law Enforcement Act (CALEA), which defines electronic surveillance guidelines for telecommunications companies.

The supported SNMPv3 traps map to messages defined by the *Lawfully Authorized Electronic Surveillance (LAES) for IP Network Access, American Nation Standard For Telecommunications*. "SNMP Traps for Subscriber Secure Policy LAES Compliance" on page 1305 describes the supported SNMPv3 traps and their related LAES messages.

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281
- SNMP Traps for Subscriber Secure Policy LAES Compliance on page 1305
- Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring on page 1307

SNMP Traps for Subscriber Secure Policy LAES Compliance

Table 118 on page 1306 describes the SNMPv3 traps that subscriber secure policy mirroring uses to provide information that maps to messages defined in the *Lawfully Authorized Electronic Surveillance (LAES) for IP Network Access, American National Standard for*

Telecommunications. These messages enable subscriber secure policy to comply with the *Communications Assistance for Law Enforcement Act* (CALEA). The Juniper Packet Mirroring MIB, **jnx-js-packet-mirror.mib**, provides the SNMP trap.

Table 118: Subscriber Secure Policy SNMPv3 Traps for LAES Messages

SNMPv3 Trap	LAES Message	Description
jnxPacketMirrorLiSubscriberLoggedIn	• **access-attempt** (implied) • **access-session-accept** • **packet-data-session-start**	A subscriber, who is identified to have a mirrored service that is activated at login, has successfully logged in.
jnxPacketMirrorSessionLiSubscriberLogInFailed	• **access-attempt** (implied) • **access-failed** (all termination reasons except authentication-reject) • **access-reject** (termination reason is authentication-reject)	A subscriber, who is identified to have a mirrored service that is activated at login, has failed to log in.
jnxPacketMirrorInterfaceLiSubscriberLoggedOut	• **access-session-end** • **packet-data-session-end**	A subscriber, who had an active mirrored service, has logged out.
jnxPacketMirrorInterfaceLiServiceActivated	• **packet-data-session-already-established**	A mirrored session has been activated.
jnxPacketMirrorSessionLiServiceActivationFailed	–	A mirrored session for a subscriber has failed.
jnxPacketMirrorSessionLiServiceDeactivated	–	A mirrored session for an established subscriber has been deactivated.
jnxPacketMirrorMirroringFailure	–	A mirrored service request failed due to an invalid value in the request. Note: This trap is not related to LAES messages.
jnxPacketMirrorTriggerType	–	The type of trigger that caused the mirroring session to be activated.
jnxPacketMirrorCallingStationIdentifier	–	The calling station ID of the subscriber whose traffic is currently being mirrored.
jnxPacketMirrorNasIdentifier	–	The NAS ID of the session in which traffic is being mirrored.
jnxPacketMirrorTargetIPv6Address	–	The IPv6 address of the subscriber interface that is being mirrored.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Intercept-Related Events Transmitted to the Mediation Device on page 1305
- Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring on page 1307

Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring

This topic provides an overview of the SNMPv3 configuration process as it pertains to subscriber secure policy. The steps are described in detail in Chapter 7, "Configuring SNMPv3" in the *Network Management Administration Guide for Routing Devices*.

To configure SNMPv3 trap support for subscriber secure policy and to send the trap information to the mediation device:

1. Configure the MIB view.

   See *Configuring MIB Views*.

2. Configure the trap notification and trap notification filter. See the following topics:
   - *Configuring the SNMPv3 Trap Notification*
   - *Configuring the Trap Notification Filter*
3. Configure the target device. The target device is the mediation device that receives the trap information.

   See *Configuring SNMPv3 Traps on a Device Running Junos OS*.

4. Configure the SNMPv3 user, authentication method and password, and privacy method and password. See the following topics:
   - *Creating SNMPv3 Users*
   - *Configuring the SNMPv3 Authentication Type*
   - *Configuring the Encryption Type*
5. Configure user access privileges to management information.

   See *Defining Access Privileges for an SNMP Group*.

Related Documentation

- Intercept-Related Events Transmitted to the Mediation Device on page 1305
- SNMP Traps for Subscriber Secure Policy LAES Compliance on page 1305
- Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring on page 1307
- For information about SNMPv3, see the *Network Management Administration Guide for Routing Devices*

Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring

This example shows an SNMP configuration that provides SNMPv3 trap support.

Configure the SNMPv3 trap support at the **[edit snmp]** hierarchy level.

```
[edit snmp]
v3 {
  usm {
    local-engine {
      user mediation-device1 { ## Name of the mediation device
        authentication-md5 {
          authentication-key "yourAuthentictaionKey"; ## SECRET-DATA
        }
        privacy-des {
          privacy-key "YourPrivacyKey"; ## SECRET-DATA
        }
      }
    }
  }
  target-address london-1 {
    address 172.19.87.240; ## Address of the mediation device receiving the traps
    port 162;
    tag-list mediation-8;
    target-parameters tp1;
  }
  target-parameters tpi {
    parameters {
      message-processing-model v3;
      security-model usm;
      security-level authentication;
      security-name mediation-device1; ## Name of the mediation device
    }
    notify-filter nf1;
  }
  notify n1 {
    type trap;
    tag mediation-8;
  }
  notify-filter nf1 {
    oid .1 include;
  }
}
view system {
  oid 1.3.6.1.2.1.1 include;
}
view all {
  oid .1 include;
}
```

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring on page 1307
- For information about SNMPv3, see the *Network Management Administration Guide for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

PART 17

Protocols for Subscriber Access

- ANCP Overview on page 1311
- Configuring ANCP on page 1331
- Dynamic IGMP Configuration Overview on page 1371
- Dynamic MLD Configuration Overview on page 1373
- Dynamic Router Advertisement Overview on page 1375

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 79

ANCP Overview

- ANCP Topology Discovery and Traffic Reporting Overview on page 1311
- Traffic Rate Reporting and Adjustment by ANCP on page 1314
- ANCP Interactions with AAA on page 1317
- ANCP Operations in Different Network Configurations on page 1319
- ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes on page 1327
- Preservation of CoS Shaping Across ANCP Restarts on page 1329

ANCP Topology Discovery and Traffic Reporting Overview

This topic describes Access Node Control Protocol (ANCP) and the ANCP agent running in Junos OS as a means to monitor subscriber access lines and to report and modify subscriber traffic on the access lines between the subscribers and the access nodes.

- Overview on page 1311
- Topology Discovery on page 1312
- Subscriber Services on page 1312
- ANCP Interfaces and Access Loop Circuit Identifiers on page 1313
- ANCP Neighbors on page 1313

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

Copyright © 2013, Juniper Networks, Inc.

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 events and operations by including the **traceoptions** statement at the **[edit protocols ancp]** hierarchy level.

Junos OS supports the following interface types for ANCP:

- Static VLAN and VLAN demux interfaces and interface sets
- Dynamic VLAN and VLAN demux interfaces and interface sets
- Dynamic VLAN-tagged interface sets
- Dynamic agent circuit identifier (ACI) interface sets
- Dynamic DHCP IP demux subscriber interfaces
- Dynamic PPPoE subscriber interfaces

ANCP was developed as an extension of *RFC 3292, General Switch Management Protocol (GSMP) V3*, but is now defined in *RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks*.

Topology Discovery

The router uses topology discovery to collect information from the access node. The information includes the following:

- Topology of the access network
- DSL line state
- Actual upstream and downstream net data rates of a synchronized DSL link
- Maximum attainable upstream and downstream net data rates
- Interleaving delay

Subscriber Services

The router receives the service profile for the subscribers from a RADIUS server. Most of the services are enforced by the router itself. The router shapes the aggregate egress traffic to subscribers based on the local loop throughput reported by the DSLAM. This traffic shaping optimizes traffic flow while avoiding traffic drops in the access node.

Some service attributes, such as interleaving delay and multicast channel information, are enforced at the access node. The ANCP agent provides the line configuration mechanism that the edge device can use to pass the line configuration to the access nodes. Typically, multiple profiles are provisioned on the access node. The router instructs the access node which profile to use for a given subscriber.

Subscribers typically receive some combination of voice, data, and video services. Each service can be provisioned on a VLAN. A subscriber might receive only a single service

 Copyright © 2013, Juniper Networks, Inc.

over a single VLAN configured on a logical interface. A group of VLANs carrying services to a subscriber is an *interface set*.

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 DSL Forum Agent-Circuit-ID VSA [26–1].

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.

Both static and dynamic ACIs must be unique across the network; no two interfaces across multiple neighbors (access nodes) can share the same identifier. ANCP support for RADIUS authentication and accounting requires that the ACIs be unique across the network. 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.

ANCP Neighbors

The ANCP agent can report traffic only for access nodes that are configured as ANCP neighbors. 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.

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 ANCP 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 ANCP to operate in a backward-compatible mode for this neighbor; for neighbors that use the original IETF implementation of ANCP 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.

Related Documentation

- Traffic Rate Reporting and Adjustment by ANCP on page 1314
- Configuring ANCP on page 1332
- Triggering ANCP OAM on page 1342
- Agent Circuit Identifier-Based Dynamic VLANs Overview on page 702

Traffic Rate Reporting and Adjustment by ANCP

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 1315
- Traffic Rate Adjustment on page 1315
- Recommended Traffic Shaping Rates on page 1316
- ANCP Keepalives on page 1316

Copyright © 2013, Juniper Networks, Inc.

Overview

ANCP 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.
- However, each DSL line type has a certain technology overhead; so 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 by a fixed percentage for each line type to generate the adjusted rate.

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 unadjusted 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. It never reports adjusted traffic rates to CoS. In addition to the unadjusted downstream rate, the agent also reports to CoS the overhead mode and bytes for the access line; CoS can use this information when it subsequently shapes the traffic.

When you remove a shaping rate configuration that the ANCP agent previously applied, the traffic shaping rate reverts to the CoS session shaping as determined by the CoS traffic control profiles specified in the dynamic profile. If ANCP remains running but loses a connection to a particular neighbor whose subscriber traffic is adjusted as a result of ANCP, 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.

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. However, when the ANCP agent subsequently reports a net data rate, by default it includes this overhead and therefore reports 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 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 VDS2. 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. That is, you configure the statements for all relevant line types, and the ANCP agent applies the appropriate adjustment when it identifies the line

Copyright © 2013, Juniper Networks, Inc.

type for the interface. The adjustment factor applies globally for all subscribers of the particular DSL line type associated with the statement: ADSL, ADSL2, ADSL+, SDS1, VDS1, or VDS2. 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, meaning that 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. However, the ANCP agent only reports net data rates to CoS; it never reports adjusted data rates to CoS. 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 ANCP-reported net data rate reduces the accuracy of CoS traffic shaping.

Recommended Traffic Shaping Rates

To handle a situation where the router does not receive information from the access node about the downstream and upstream calculated traffic rates for an interface, you can specify recommended *advisory* values for shaping the traffic sent to the interface so that it matches the subscriber local loop speed.

The transmit speed is the recommended traffic value in bits per second used for downstream traffic for an ANCP interface, and is conveyed in the Juniper Networks VSA, Downstream-Calculated-Qos-Rate (IANA 4874, 26–141). The receive speed is the recommended traffic value in bits per second used for upstream traffic for an ANCP interface, and is conveyed in the Juniper Networks VSA, Upstream-Calculated-Qos-Rate VSA (IANA 4874, 26-142).

To set the recommended shaping rates that are used as the default values for these VSAs in static configurations, include the **downstream-rate** and **upstream-rate** statements at the **[edit interfaces *interface-name* unit *logical-unit-number* advisory-options]** hierarchy level.

To configure the recommended rates on dynamically created VLAN interfaces, include the **upstream-rate** or **downstream-rate** statements at the **[edit dynamic-profiles *profile-name* interfaces $junos-interface-ifd-name unit *$junos-interface-unit* advisory-options]** hierarchy level.

To configure the recommended rates on dynamically created ACI interface sets, include the **upstream-rate** or **downstream-rate** statements at the **[edit dynamic-profiles *profile-name* interface-set $junos-interface-set-name interfaces $junos-interface-ifd-name advisory-options]** hierarchy level.

ANCP Keepalives

ANCP 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 ANCP. 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

 Copyright © 2013, Juniper Networks, Inc.

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 ANCP sends keepalive messages every 40 seconds. In this example, if CoS does not receive a keepalive message within 120 seconds, then it reverts any ANCP-derived policy changes.

Related Documentation

- ANCP Topology Discovery and Traffic Reporting Overview on page 1311
- Configuring ANCP on page 1332
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 1082
- Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 1083
- ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes on page 1327

ANCP 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 access line from the ANCP 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. 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

Copyright © 2013, Juniper Networks, Inc.

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 CoS configuration. The agent then maps the adjusted rates to Upstream-Calculated-Qos-Rate [26-142] and Downstream-Calculated-Qos-Rate [26-141].
2. If the session cannot be correlated with an access line, but the PPPoE or DHCP discovery packet includes the DSL Forum VSA and the Access-Loop-Encapsulation subattribute includes a value for the AAL5 data link, then the ANCP agent adjusts the Actual-Data-Rate-Downstream and Actual-Data-Rate-Upstream subattributes to account for the ATM 48/53 cell tax. The adjusted rates mapped to Upstream-Calculated-Qos-Rate [26-142] and Downstream-Calculated-Qos-Rate [26-141].
3. If neither of the preceding sets of conditions is satisfied, then the ANCP agent simply maps the recommended downstream and upstream data rates to Upstream-Calculated-Qos-Rate [26-142] and Downstream-Calculated-Qos-Rate [26-141]. The recommended rates are either configured statically for the VLAN or VLAN demux interfaces or are in the dynamic profile that crates the interfaces.

To map an ACI to a static VLAN demux interface, include the **access-identifier** ***identifier*** statement at the **[edit protocols ancp interfaces demux0.*logical-unit-number*]** hierarchy level.

To configure advisory upstream and downstream data rates on a static VLAN demux interface, include the **upstream-rate** ***rate*** or **downstream-rate** ***rate*** statements at the **[edit interfaces demux0 unit *logical-unit-number*]** hierarchy level.

To configure an underlying interface for the PPPoE sessions in an interface set, include the **underlying-interface** ***interface-name*** statement at the **[edit protocols ancp interfaces interface-set *interface-set-name*]** hierarchy level.

When an ACI, and therefore a subscriber access line, has been mapped to an interface or interface set, the ACI can be re-mapped to a different interface or set. When this happens, traffic shaping is adjusted accordingly for the interfaces or interface sets involved. This capability is useful for the Business Services model, where a PPPoE session that is initially classified as a residential household can be reclassified as a business subscriber during RADIUS authentication by using a Junos OS ICE AAA framework Op-Script application.

In the Business Services Model, the PPPoE session initially represents a residential household until RADIUS authentication and authorization takes place. The ANCP agent

 Copyright © 2013, Juniper Networks, Inc.

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 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 Topology Discovery and Traffic Reporting Overview on page 1311
- ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes on page 1327
- Configuring ANCP on page 1332
- Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages on page 1341

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 ANCP for any of the following interface types:

- Static VLAN (over Ethernet) and VLAN demux (over aggregated Ethernet) interfaces and interface sets.
- Dynamic VLAN and VLAN demux interfaces and interface sets
- Dynamic VLAN-tagged interface sets
- Dynamic ACI interface sets
- Dynamic DHCP IP demux subscriber interfaces
- Dynamic PPPoE subscriber interfaces

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.

The 1:1 and N:1 configuration models determine how VLANs are correlated with households. 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

Copyright © 2013, Juniper Networks, Inc.

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 119 on page 1320 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 119: 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 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 120 on page 1320 describes the types of interface sets supported in the ANCP 1:1 access model, how they are created, and how the ACI is mapped to the interface set.

Table 120: 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.

 Copyright © 2013, Juniper Networks, Inc.

Table 120: ACI Mapping by Interface Set Type for the ANCP N:1 Access Model *(continued)*

Type of Interface Set	Description	Interface Type	Presence of ACI in Discovery Packets
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 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 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 configuration must include the **access-identifier** statement to statically map the ACI to the interface set. All DHCP IP demux or PPPoE sessions that have the same VLAN ID tag are mapped to the same interface set.	Dynamic VLAN and VLAN demux interfaces.	Irrelevant.

CoS traffic shaping is based on the subscriber downstream traffic rate that the ANCP agent receives from the access node and then passes to CoS. CoS can shape subscriber traffic at the level of the household or the session:

- Household shaping—Only aggregate traffic to the household is shaped. Household shaping results from applying a CoS traffic control profile to the static VLAN or VLAN demux interface or to the interface set.
- Session shaping—The traffic rate to individual devices in the household is shaped. Session shaping results from specifying a CoS traffic control profile in the dynamic PPPoE profile that creates the subscriber session. Depending on the network configuration, session shaping may employ shared priority queues to shape all sessions identically or individual priority queues to shape the sessions separately.

The following sections illustrate several possible configurations and lists the sequence of events for the ANCP operations in each case. Not every possible configuration is presented.

Copyright © 2013, 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 49 on page 1322). 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 49: Sample ANCP Topology Without Interface Sets (1:1 and N:1 Model)

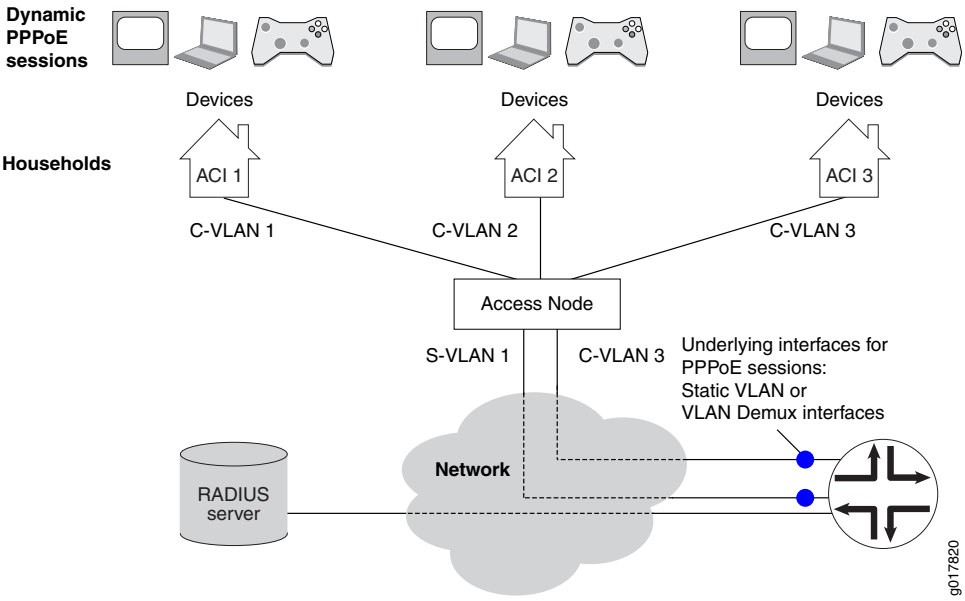

Copyright © 2013, Juniper Networks, Inc.

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 49 on page 1322 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 © 2013, 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 50 on page 1324). 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 50: Sample ANCP Topology with Interface Sets (N:1 Model)

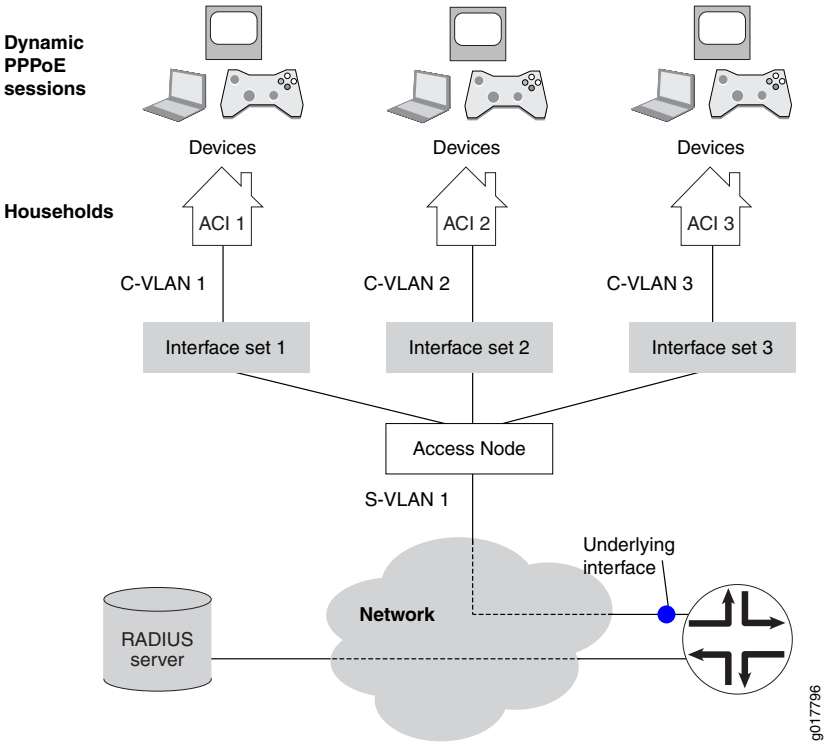

In this N:1 model with interface sets, the access node must add the DSL Forum VSA to the PPPoE PADI and PADR discovery packets that it passes to the router during the establishment of dynamic PPPoE sessions. The VSA includes the ACI for the household. This inclusion enables AAA to correlate the PPPoE 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 © 2013, 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 50 on page 1324 with static VLAN interfaces over Ethernet with interface sets.

1. A network device in the household initiates PPPoE discovery.
2. The access node adds the DSL Forum VSA tag with the ACI for the household to the PPPoE PADI and PADR discovery packets. (The identifier is known to PPPoE as the agent circuit identifier.)
3. PPPoE creates a dynamic PPPoE session with the provided ACI on the underlying static VLAN and applies the advisory options configured on the VLAN to the session.
4. The access node independently provides the ANCP agent with the ANCP DSL attributes for an access line identified by an ACI.
5. The ANCP agent provides CoS with the adjusted downstream data rate for the interface set mapped to the ACI. The ANCP agent stores all ANCP DSL attributes, including the adjusted upstream and downstream data rates, in the router's shared database.
6. AAA correlates the dynamic PPPoE session with the access line by matching the session identifier received in the DSL Forum VSA to the ACI configured for the interface set in the ANCP agent configuration.
7. AAA retrieves the ANCP DSL attributes for the access line from the router's shared database and maps them to the Juniper Networks DSL VSAs in the RADIUS Access-Request and Accounting-Request messages. If the DSL attributes are unavailable, the session's advisory upstream and downstream data rates are mapped to the Upstream-Calculated-Qos-Rate VSA (26-142) and Downstream-Calculated-Qos-Rate (26–141) VSAs, respectively. These VSAs are then included in the RADIUS messages.
8. When authentication is completed, the dynamic PPPoE session is placed into the interface set configured in the dynamic PPPoE profile. The profile specifies a named interface set or the **$junos-interface-set-name** predefined variable, which indicates that the interface set is named in the RADIUS Access-Accept message.

Copyright © 2013, Juniper Networks, Inc.

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 51 on page 1326). 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 51: Sample ANCP Topology with Interface Sets (1:1 Model)

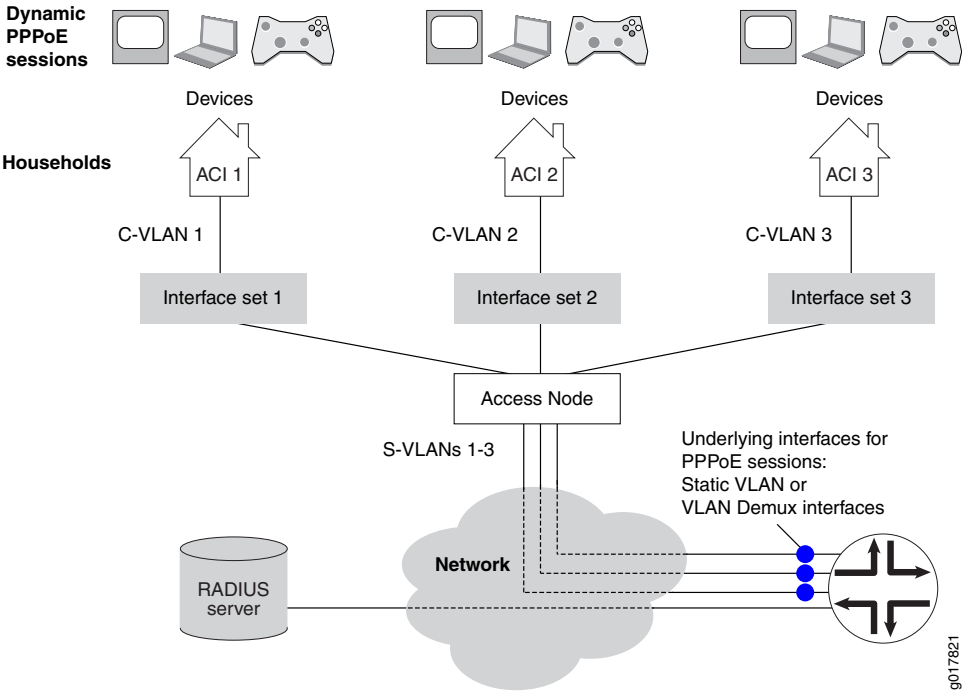

In this 1:1 model with interface sets, the ANCP agent configuration must map the underlying interface for the PPPoE sessions in an interface set to both the ACI and the interface set. This configuration enables AAA to correlate the PPPoE sessions with their respective subscriber access lines and DSL attributes during RADIUS authentication and accounting.

When the dynamic PPPoE profile is configured with the **$junos-interface-set-name** predefined variable, the configuration of the access node, router, and RADIUS server must be synchronized with regard to the ACI and interface set:

- The RADIUS Access-Accept message must contain the Juniper Networks Qos-Interface-Set-Name VSA (26-130).
- The CoS Layer 2 configuration must explicitly identify the interface set that is named in the Qos-Interface-Set-Name VSA (26-130).
- The ANCP agent configuration must map an ACI to the interface set that is named in the Qos-Interface-Set-Name VSA (26-130).

Copyright © 2013, Juniper Networks, Inc.

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 51 on page 1326 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 Topology Discovery and Traffic Reporting Overview on page 1311
- Configuring ANCP on page 1332
- Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet on page 1349

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.

Copyright © 2013, Juniper Networks, Inc.

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.

Table 121 on page 1328 shows the relationship between the ANCP DSL attributes, Juniper Networks DSL VSAs, and DSL Forum VSAs.

Table 121: 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 [26–115]—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 [26–114]—Unadjusted rate • Upstream-Calculated-Qos-Rate [26–142]—Rate as adjusted by ANCP	Actual-Data-Rate-Upstream [26–129]
Attainable-Net-Data-Rate-Upstream	Att-Data-Rate-Up [26–117]	Attainable-Data-Rate-Upstream [26–133]

Copyright © 2013, Juniper Networks, Inc.

Table 121: 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]
Attainable-Net-Data-Rate-Downstream	Att-Data-Rate-Dn [26–118]	Attainable-Data-Rate-Downstream [26–134]
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]
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 98
- Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages on page 1341

Preservation of CoS Shaping Across ANCP Restarts

When the ANCP agent stops due to a process or GRES, CoS enforces the ANCP downstream shaping-rates until the CoS keepalive timer expires. When the timer expires, CoS reverts to the CoS shaping-rate configured for the interfaces.

You configure the CoS keepalive timer by including the **maximum-helper-restart-time *seconds*** statement at the **[edit protocols ancp]** hierarchy level. It specifies how much time other daemons such as CoS wait for ANCP 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 Topology Discovery and Traffic Reporting Overview on page 1311
- Configuring ANCP on page 1332
- Specifying How Long Processes Wait for ANCP Restart to Complete on page 1337

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 80

Configuring ANCP

- Configuring ANCP on page 1332
- Configuring ANCP Neighbors on page 1333
- Associating an Access Node with Subscribers for ANCP Operations on page 1334
- Specifying the Interval Between ANCP Adjacency Messages on page 1335
- Specifying the Maximum Number of Discovery Table Entries on page 1335
- Configuring ANCP for Backward Compatibility on page 1336
- Specifying How Long Processes Wait for ANCP Restart to Complete on page 1337
- Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates on page 1337
- Configuring the ANCP Agent to Report Traffic Rates to CoS on page 1338
- Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces on page 1339
- Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages on page 1341
- Triggering ANCP OAM on page 1342
- Verifying and Monitoring ANCP Neighbors on page 1343
- Clearing ANCP Neighbors on page 1343
- Verifying and Monitoring ANCP Subscribers on page 1344
- Clearing ANCP Subscribers on page 1344
- Verifying and Monitoring CoS for ANCP Subscribers on page 1345
- Tracing ANCP Operations for Subscriber Access on page 1345
- Configuring the ANCP Trace Log Filename on page 1346
- Configuring the Number and Size of ANCP Log Files on page 1346
- Configuring Access to the ANCP Log File on page 1347
- Configuring a Regular Expression for ANCP Messages to Be Logged on page 1347
- Configuring the ANCP Tracing Flags on page 1348
- Configuring the Severity Level to Filter Which ANCP Messages Are Logged on page 1348
- Example: Configuring an ANCP Network with Interface Sets and N:1 Static Demux VLANs over Aggregated Ethernet on page 1349

Copyright © 2013, Juniper Networks, Inc.

Configuring ANCP

You can configure ANCP to enable a service-oriented Layer 3 edge device to discover information about the topology of a connected access network. ANCP can also provide details about subscriber traffic and enable the adjustment of QoS traffic shaping for subscribers.

To configure ANCP:

1. Specify each ANCP neighboring access node to be monitored and optionally configure neighbor parameters.

   See "Configuring ANCP Neighbors" on page 1333.

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 Operations" on page 1334.

3. (Optional) Configure the adjacency timer.

   See "Specifying the Interval Between ANCP Adjacency Messages" on page 1335.

4. (Optional) Specify the maximum number of discovery table entries that are accepted.

   See "Specifying the Maximum Number of Discovery Table Entries" on page 1335

5. (Optional) Configure ANCP to work with an early IETF draft.

   See "Configuring ANCP for Backward Compatibility" on page 1336.

6. (Optional) Configure the graceful restart timer.

   See "Specifying How Long Processes Wait for ANCP Restart to Complete" on page 1337.

7. (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 1337.

8. (Optional) Configure the ANCP agent to report unadjusted downstream traffic rates to CoS.

   See "Configuring the ANCP Agent to Report Traffic Rates to CoS" on page 1338.

9. (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 1339.

10. (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 1341.

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

 Copyright © 2013, Juniper Networks, Inc.

See "Tracing ANCP Operations for Subscriber Access" on page 1345.

Related Documentation

- ANCP Topology Discovery and Traffic Reporting Overview on page 1311
- Triggering ANCP OAM on page 1342

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 10.2.3.4
   ```

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 10.2.3.4]
   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 10.2.3.4]
   user@host# set ietf-mode
   ```

4. (Optional) Configure the interval in seconds between ANCP adjacency messages exchanged with this neighbor.

   ```
   [edit protocols ancp neighbor 10.2.3.4]
   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 10.2.3.4]
   user@host# set maximum-discovery-table-entries 10000
   ```

Related Documentation

- Configuring ANCP on page 1332
- Configuring ANCP for Backward Compatibility on page 1336
- Specifying the Interval Between ANCP Adjacency Messages on page 1335
- Specifying the Maximum Number of Discovery Table Entries on page 1335

Copyright © 2013, Juniper Networks, Inc.

Associating an Access Node with Subscribers for ANCP 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.

NOTE: In earlier Junos OS releases, we recommend that the ACIs be unique across all neighbors, but this is not required. When the ACIs are not unique in the earlier releases, you can configure the subscribers to be uniquely identified by associating the IP address of the neighbor with the subscriber's interface or interface-set. This configuration is no longer supported and the ACIs must be unique across your ANCP network.

The ACIs can be statically or dynamically configured. 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.

To associate an ACI with a set of VLAN interfaces for subscribers:

- Specify the name of the interface set and the unique ACI for the access node.

```
[edit protocols ancp interfaces]
user@host# set interface-set vlan5 access-identifier "dslam port 2/3"
```

To associate an ACI with a single VLAN:

- Specify the logical interface and the unique ACI for the access node.

```
[edit protocols ancp interfaces]
user@host# set ge-1/0/4.12 access-identifier "dslam port-2-10"
```

To associate an ACI with a static VLAN demux interface:

- Specify the logical interface and the unique ACI for the access node.

```
[edit protocols ancp interfaces]
user@host# set demux0.100 access-identifier aci_100_1_0
```

Related Documentation

- Configuring ANCP on page 1332
- interfaces on page 1732

Copyright © 2013, Juniper Networks, Inc.

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 10.2.3.4]
user@host# set adjacency-timer 20
```

Related Documentation

- Configuring ANCP on page 1332
- Configuring ANCP Neighbors on page 1333

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 10.2.3.4]
user@host# set maximum-discovery-table-entries 5000
```

Related Documentation

- Configuring ANCP on page 1332
- Configuring ANCP Neighbors on page 1333

Copyright © 2013, Juniper Networks, Inc.

Configuring ANCP for Backward Compatibility

You can configure ANCP 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 10.2.3.4]
  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 10.2.3.4]
  user@host# set ietf-mode
  ```

Related Documentation

- Configuring ANCP on page 1332
- Configuring ANCP Neighbors on page 1333

Copyright © 2013, Juniper Networks, Inc.

Specifying How Long Processes Wait for ANCP 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 ANCP to be down and immediately reverts any traffic shaping updates that were implemented as a result of ANCP 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 ANCP comes back up and sends fresh traffic shaping updates to CoS.

To configure how long other processes wait for ANCP to restart:

- Specify the time in seconds.

```
[edit protocols ancp]
user@host# set maximum-helper-restart-time 150
```

Related Documentation

- Configuring ANCP on page 1332
- *Configuring ANCP to Adjust CoS Traffic Shaping*
- **qos-adjust on page 1910**

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 therefore 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. That is, you configure the statements for all relevant line types, and the agent applies the appropriate adjustment when it identifies the line type for the interface. 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.

Copyright © 2013, Juniper Networks, Inc.

NOTE: The ANCP agent reports only unadjusted downstream data rates to CoS, and reports 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:

- 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-sds1 adjustment-factor
user@host# set qos-adjust-vds1 adjustment-factor
user@host# setqos-adjust-vds2 adjustment-factor
```

Related Documentation

- Configuring ANCP on page 1332
- ANCP Topology Discovery and Traffic Reporting Overview on page 1311

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 unadjusted or net downstream data rates to CoS for all subscribers in the network. This information enables CoS to subsequently shape the traffic on these 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 therefore 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, nor does it report adjusted downstream traffic rates to CoS.

If CoS does not receive a keepalive message within the maximum helper restart time, it considers ANCP to be down and immediately reverts to the configured values any traffic shaping updates that were modified as a result of traffic reports from the ANCP agent. The configured values are maintained until ANCP comes back up and sends fresh traffic updates to CoS.

However, if ANCP 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.

 Copyright © 2013, Juniper Networks, Inc.

To configure the ANCP agent to report unadjusted downstream traffic rates to CoS for traffic shaping:

- Specify that the ANCP agent reports traffic rates to CoS.

```
[edit protocols ancp]
user@host# set qos-adjust
```

Related Documentation

- Traffic Rate Reporting and Adjustment by ANCP on page 1314
- Configuring ANCP on page 1332
- ANCP Topology Discovery and Traffic Reporting Overview on page 1311
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 1082
- Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 1083
- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1099
- Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 1104
- Specifying How Long Processes Wait for ANCP Restart to Complete on page 1337
- **maximum-helper-restart-time on page 1777**

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.

```
[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
```

Related Documentation

- Configuring ANCP on page 1332
- Configuring AAA to Include Juniper Networks DSL VSAs in RADIUS Messages on page 1341

Copyright © 2013, Juniper Networks, Inc.

- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- ANCP Topology Discovery and Traffic Reporting Overview on page 1311

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 1327 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 108 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.

  [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** exclude**max-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 ANCP on page 1332
- Setting a Recommended Shaping Rate for Traffic on ANCP Interfaces on page 1339

Copyright © 2013, Juniper Networks, Inc.

- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes on page 1327
- ANCP Topology Discovery and Traffic Reporting Overview on page 1311

Triggering ANCP OAM

You can trigger ANCP OAM to perform a loopback test on the local loop between the access node and the CPE to help isolate simple faults. On an ATM-based local loop, the ANCP operation triggers the access node to generate ATM (F4/F5) loopback cells on the local loop. On an Ethernet-based local loop, the ANCP operation triggers the access node to generate an Ethernet loopback message on the local loop. When the test completes, the access node sends a message to the router with the results.

Issue the **request ancp oam neighbor** command from CLI operational mode to initiate testing of a local loop identified by the IP address or system name of the ANCP neighbor and the ACI for a subscriber on that access node.

Issue the **request ancp oam interface** command from CLI operational mode to initiate testing of a local loop identified by the ANCP interface or interface set associated with a subscriber and the ACI for a subscriber on that access node.

With both commands, you can also specify how many times the test must be run and how long the router waits for a response to the OAM request.

To initiate ANCP local loop testing:

- Identify the loop by the subscriber identifier and the neighbor's IP address; optionally specify how many times the test runs and the timeout period.

  user@host> **request ancp oam neighbor ip-address 192.168.32.5 subscriber "dslam port-2-10" count 5 timeout 600**

- Identify the loop by the subscriber identifier and the neighbor's system name; optionally specify how many times the test runs and the timeout period.

  user@host> **request ancp oam neighbor system-name ba:ad:be:ef:10:10 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 Topology Discovery and Traffic Reporting Overview on page 1311
- Configuring ANCP on page 1332

 Copyright © 2013, Juniper Networks, Inc.

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 10.25.64.21**

- To display detailed information, add **detail** to the command:

  user@host> **show ancp neighbor detail**

  user@host> **show ancp neighbor ba:ad:be:ef:10:10 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 10.25.64.21**

- To display total and state-wise counts of both ANCP neighbors and subscribers:

  user@host> **show ancp summary**

Related Documentation

- *Junos OS Operational Mode Commands*

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 10.25.64.21**

  user@host> **clear ancp neighbor system-name ba:ad:be:ef:10:10**

- To verify that the connection has been cleared:

  user@host> **show ancp neighbor**

  user@host> **show ancp neighbor 10.25.64.21**

  user@host> **show ancp neighbor ba:ad:be:ef:10:10**

Related Documentation

- *Junos OS Operational Mode Commands*

Copyright © 2013, Juniper Networks, Inc.

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 10.25.64.21**

- 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 10.25.64.21 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**

Related Documentation

- *Junos OS Operational Mode Commands*

Clearing ANCP Subscribers

Purpose Clear ANCP subscriber information.

Action

- To clear connections with all ANCP subscribers:

  user@host> **clear ancp subscriber**

- 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 10.25.64.21**

  user@host> **clear ancp subscriber identifier port-2-10 system-name ba:ad:be:ef:10:10**

- To verify that the connection has been cleared:

  user@host> **show ancp subscriber**

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *Junos OS Operational Mode Commands*

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**

Related Documentation

- *Junos OS Operational Mode Commands*

Tracing ANCP Operations for Subscriber Access

The Junos OS trace feature tracks ANCP 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 *Junos OS System Log Messages Reference*.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure ANCP tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the ANCP Trace Log Filename" on page 1346.

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

   See "Configuring the Number and Size of ANCP Log Files" on page 1346.

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

   See "Configuring Access to the ANCP Log File" on page 1347.

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 1347.

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

   See "Configuring the ANCP Tracing Flags" on page 1348.

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 1348.

Related Documentation

- Configuring ANCP on page 1332

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 Operations for Subscriber Access on page 1345

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

 Copyright © 2013, Juniper Networks, Inc.

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 Operations for Subscriber Access on page 1345

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 Operations for Subscriber Access on page 1345

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 Operations for Subscriber Access on page 1345

Configuring the ANCP Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

  [edit protocols ancp traceoptions]
  user@host# **set flag restart**

Related Documentation

- Tracing ANCP Operations for Subscriber Access on page 1345

Configuring the Severity Level to Filter Which ANCP Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all** or **verbose**. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages . By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

  [edit protocols ancp traceoptions]
  user@host# **set level** ***severity***

Related Documentation

- Tracing ANCP Operations for Subscriber Access on page 1345
- Configuring ANCP on page 1332
- **traceoptions on page 2033**

 Copyright © 2013, Juniper Networks, Inc.

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 1349
- Overview on page 1349
- Configuration on page 1354
- Verification on page 1366

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
- Junos OS Release 12.2 or later

Before you begin configuring the example, be sure you have:

- Thoroughly read and understood the following topics:
  - ANCP Topology Discovery and Traffic Reporting Overview on page 1311
  - ANCP Operations in Different Network Configurations on page 1319
- 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.

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 52 on page 1350 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 52: N:1 ANCP Topology with Interface Sets and VLAN Demux Interface over Aggregated Ethernet

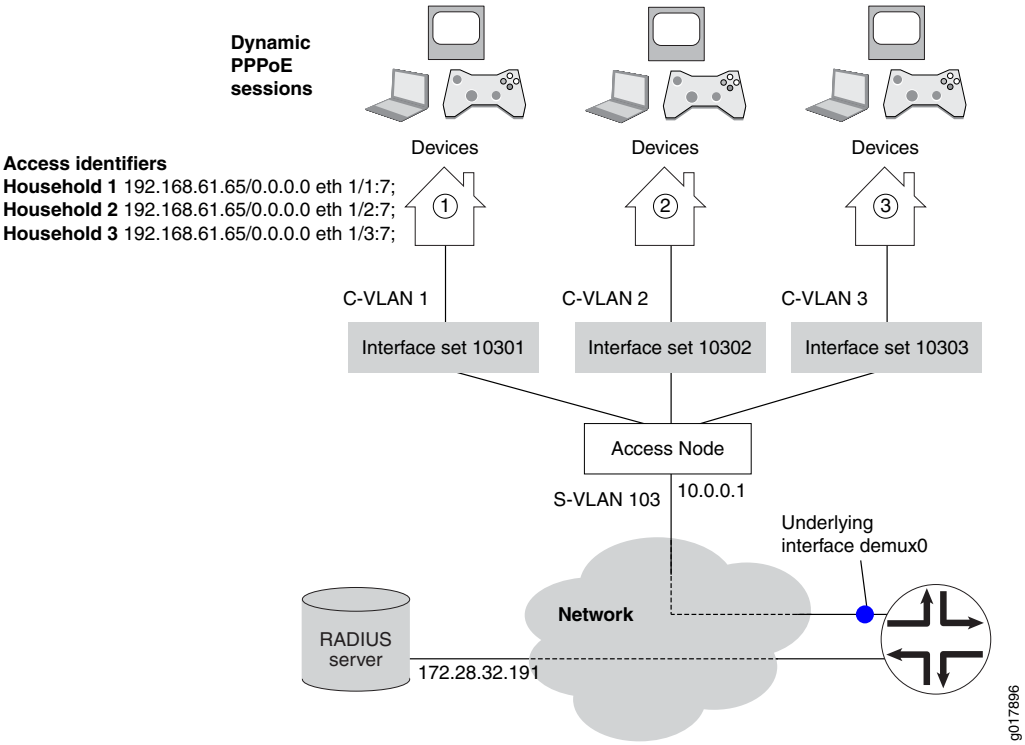

Table 122 on page 1351 describes the configuration components used in this example.

Copyright © 2013, Juniper Networks, Inc.

Table 122: Configuration Components used in ANCP N:1 Topology Example with Interface Sets

Configuration Component or Property	Component Name or Setting	Description
Dynamic profiles	ancp-10301 ancp-10302 ancp-10303	Each profile defines the dynamic PPPoE session created when any of the devices for a particular subscriber household accesses the network. Each profile specifies the following: • A set of interfaces in which the sessions are created. • Dynamic instantiation of both the logical interfaces for the sessions and the underlying PPPoE logical interfaces on which the subscribers log in. • CHAP and PAP authentication for the sessions. • The interval between successive PPP keepalive messages. • The loopback address for the dynamic PPPoE logical interfaces.
Predefined variables	$junos-interface-unit	Instantiates the logical interface for each PPPoE session.
	$junos-underlying-interface	Instantiates the logical underlying PPP interface on which each dynamic PPPoE logical interface is created when a subscriber logs in.

Table 122: 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 © 2013, Juniper Networks, Inc.

Table 122: 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	10.0.0.1	Address of the ANCP access node that monitors the subscriber households.
	100.0.0.1/32	Address of the loopback interface, lo0.
	172.28.32.191	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 10.0.0.1. 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 © 2013, Juniper Networks, Inc.

The RADIUS configuration on the router includes the following elements:

- The IP address (172.28.32.191) 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 1356
- Configuring the Static VLAN Demux Interface over Aggregated Ethernet on page 1358
- Configuring Class of Service on page 1362
- Configuring ANCP on page 1363
- Configuring RADIUS Authentication and Accounting on page 1364

CLI Quick Configuration

To quickly configure the ANCP network described in this example, copy the following commands, paste them in a text file, remove any line breaks, and then copy and paste the commands into the CLI.

```
# Dynamic Profiles
edit dynamic-profiles ancp-10301
set interfaces interface-set 10301 interface pp0 unit "$junos-interface-unit"
edit interfaces pp0 unit "$junos-interface-unit"
set ppp-options chap
set ppp-options pap
set pppoe-options underlying-interface "$junos-underlying-interface"
set keepalives interval 30
set family inet unnumbered-address lo0.0
top
edit dynamic-profiles ancp-10302
set interfaces interface-set 10302 interface pp0 unit "$junos-interface-unit"
edit interfaces pp0 unit "$junos-interface-unit"
set ppp-options chap
set ppp-options pap
set pppoe-options underlying-interface "$junos-underlying-interface"
set keepalives interval 30
set family inet unnumbered-address lo0.0
top
edit dynamic-profiles ancp-10303
set interfaces interface-set 10303 interface pp0 unit "$junos-interface-unit"
edit interfaces pp0 unit "$junos-interface-unit"
```

 Copyright © 2013, Juniper Networks, Inc.

```
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 100.0.0.1/32
top
#
# Class of Service
edit class-of-service
set traffic-control-profiles tcp1 shaping-rate 16m
set traffic-control-profiles tcp1 overhead-accounting cell-mode
set interfaces interface-set 10301 output-traffic-control-profile tcp1
set interfaces interface-set 10302 output-traffic-control-profile tcp1
set interfaces interface-set 10303 output-traffic-control-profile tcp1
top
#
# ANCP
edit protocols ancp
set traceoptions file ancpd
```

```
set traceoptions file size 512m
set traceoptions flag config
set traceoptions flag cos
set qos-adjust
set adjacency-timer 5
set maximum-helper-restart-time 90
set qos-adjust-adsl 90
set qos-adjust-adsl2 95
set interfaces interface-set 10301 access-identifier "192.168.61.65/0.0.0.0 eth 1/1:7;"
set interfaces interface-set 10302 access-identifier "192.168.61.65/0.0.0.0 eth 1/2:7;"
set interfaces interface-set 10303 access-identifier "192.168.61.65/0.0.0.0 eth 1/3:7;"
set interfaces interface-set 10301 underlying-interface demux0.10301
set interfaces interface-set 10302 underlying-interface demux0.10302
set interfaces interface-set 10303 underlying-interface demux0.10303
set neighbor 10.0.0.1
top
#
# RADIUS
edit access
set radius-server 172.28.32.191 secret "$9$MUeL7VgoGqmTwYmTz3tpWLx"
edit access profile radius-profile
set authentication-order radius
set radius authentication-server 172.28.32.191
set radius accounting-server 172.28.32.191
set radius options accounting-session-id-format decimal
set radius options juniper-dsl-attributes
set radius attributes exclude dsl-forum-attributes access-request
set radius attributes exclude dsl-forum-attributes accounting-start
set radius attributes exclude dsl-forum-attributes accounting-stop
top
```

Configuring the Dynamic PPPoE Profiles

Step-by-Step Procedure

In this procedure, you configure a dynamic profile for each C-VLAN: ancp-10301, ancp-10302, and ancp-1033.

1. Configure the interface set that the PPPoE sessions on this C-VLAN are placed in.

   ```
   [edit dynamic-profiles ancp-10301]
   user@host1# edit interfaces interface-set 10301
   ```

2. Configure the logical interfaces to be dynamically instantiated for the interface set.

   ```
   [edit dynamic-profiles ancp-10301 interfaces interface-set 10301]
   user@host1# set interface pp0 unit “$junos-interface-unit”
   ```

3. Configure CHAP and PAP authentication as properties of the dynamic PPPoE logical interfaces.

   ```
   [edit dynamic-profiles ancp-10301 interfaces pp0 unit “$junos-interface-unit”]
   user@host1# set ppp-options chap
   user@host1# set ppp-options pap
   ```

4. Configure the logical underlying interface on which the router creates the dynamic PPPoE logical interface; this is the interface on which the subscriber logs in.

   ```
   [edit dynamic-profiles ancp-10301 interfaces pp0 unit “$junos-interface-unit”]
   user@host1# set pppoe-options underlying-interface “$junos-underlying-interface”
   ```


Copyright © 2013, Juniper Networks, Inc.

5. Specify the interval between successive keepalive requests.

   [edit dynamic-profiles ancp-10301 interfaces pp0 unit “$junos-interface-unit”]
   user@host1# **set keepalives interval 30**

6. Configure the IPv4 protocol family and that the local (unnumbered) address can be derived from the loopback address for the dynamic PPPoE logical interfaces.

   [edit dynamic-profiles ancp-10301 interfaces pp0 unit “$junos-interface-unit”]
   user@host1# **set family inet unnumbered-address lo0.0**

7. Repeat Steps 1 through 6 for the second dynamic profile, ancp-10302, and the third dynamic profile, ancp-10303.

Results From configuration mode, confirm the dynamic profile configuration by entering the **show dynamic-profiles** command.

```
[edit]
user@host# show dynamic-profiles
ancp-10301 {
   interfaces {
      interface-set 10301 {
         interface pp0 {
            unit "$junos-interface-unit";
         }
      }
      pp0 {
         unit "$junos-interface-unit" {
            ppp-options {
               chap;
               pap;
            }
            pppoe-options {
               underlying-interface "$junos-underlying-interface";
            }
            keepalives interval 30;
            family inet {
               unnumbered-address lo0.0;
            }
         }
      }
   }
}
ancp-10302 {
   interfaces {
      interface-set 10302 {
         interface pp0 {
            unit "$junos-interface-unit";
         }
      }
      pp0 {
         unit "$junos-interface-unit" {
            ppp-options {
               chap;
               pap;
            }
```

```
            pppoe-options {
              underlying-interface "$junos-underlying-interface";
            }
            keepalives interval 30;
            family inet {
              unnumbered-address lo0.0;
            }
          }
        }
      }
    }
    ancp-10303 {
      interfaces {
        interface-set 10303 {
          interface pp0 {
            unit "$junos-interface-unit";
          }
        }
        pp0 {
          unit "$junos-interface-unit" {
            ppp-options {
              chap;
              pap;
            }
            pppoe-options {
              underlying-interface "$junos-underlying-interface";
            }
            keepalives interval 30;
            family inet {
              unnumbered-address lo0.0;
            }
          }
        }
      }
    }
```

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

Configuring the Static VLAN Demux Interface over Aggregated Ethernet

Step-by-Step Procedure

1. Enable hierarchical scheduling on this interface.

   [edit interfaces ge-1/0/1]
   user@host1# **set hierarchical-scheduler**

2. Specify this interface as the primary member of the aggregated Ethernet bundle.

   [edit interfaces ge-1/0/1]
   user@host1# **set gigether-options 802.3ad ae0 primary**

3. Enable hierarchical scheduling on a second interface.

   [edit interfaces ge-1/0/2]
   user@host1# **set hierarchical-scheduler**

4. Specify this interface as the backup member of the aggregated Ethernet bundle.

   [edit interfaces ge-1/0/2]

Copyright © 2013, Juniper Networks, Inc.

```
user@host1# set gigether-options 802.3ad ae0 backup
```

5. Enable hierarchical scheduling on the aggregated Ethernet interface.

```
[edit interfaces ae0]
user@host1# set hierarchical-scheduler
```

6. Enable stacked VLAN tagging for all logical interfaces on the aggregated Ethernet interface.

```
[edit interfaces ae0]
user@host1# set stacked-vlan-tagging
```

7. Enable link protection as a property of the aggregated Ethernet interface.

```
[edit interfaces ae0]
user@host1# set aggregated-ether-options link-protection
```

8. Configure VLAN demux interface demux0.10301.

   a. Configure the router to respond to ARP requests on the interface.

```
[edit interfaces demux0 unit 10301]
user@host1# set proxy-arp
```

   b. Configure the outer VLAN tag to identify the access node (S-VLAN) and the inner VLAN tag to identify the subscriber port on the access node (C-VLAN).

```
[edit interfaces demux0 unit 10301]
user@host1# set vlan tags outer 103 inner 1
```

   c. Specify that the VLAN demux interface runs on the underlying aggregated Ethernet interface.

```
[edit interfaces demux0 unit 10301]
user@host1# set demux-options underlying-interface ae0
```

   d. Prevent multiple PPPoE sessions from being created for the same PPPoE subscriber on this VLAN demux interface.

```
[edit interfaces demux0 unit 10301]
user@host1# set family pppoe duplicate-protection
```

   e. Configure the dynamic profile that is instantiated on the VLAN demux interface.

```
[edit interfaces demux0 unit 10301]
user@host1# set family pppoe dynamic-profile ancp-10301
```

   f. Configure the recommended upstream and downstream traffic rates.

```
[edit interfaces demux0 unit 10301]
user@host1# set advisory-options upstream-rate 1m
user@host1# set advisory-options downstream-rate 16m
```

9. Configure VLAN demux interface demux0.10302.

   a. Configure the router to respond to ARP requests on the interface.

```
[edit interfaces demux0 unit 10302]
user@host1# set proxy-arp
```

   b. Configure the outer VLAN tag to identify the access node (S-VLAN) and the inner VLAN tag to identify the subscriber port on the access node (C-VLAN).

```
[edit interfaces demux0 unit 10302]
user@host1# set vlan tags outer 103 inner 2
```

c. Specify that the VLAN demux interface runs on the underlying aggregated Ethernet interface.

```
[edit interfaces demux0 unit 10302]
user@host1# set demux-options underlying-interface ae0
```

d. Prevent multiple PPPoE sessions from being created for the same PPPoE subscriber on this VLAN demux interface.

```
[edit interfaces demux0 unit 10302]
user@host1# set family pppoe duplicate-protection
```

e. Configure the dynamic profile that is instantiated on the VLAN demux interface.

```
[edit interfaces demux0 unit 10302]
user@host1# set family pppoe dynamic-profile ancp-10302
```

f. Configure the recommended upstream and downstream traffic rates.

```
[edit interfaces demux0 unit 10302]
user@host1# set advisory-options upstream-rate 1m
user@host1# set advisory-options downstream-rate 16m
```

10. Configure VLAN demux interface demux0.10303.

a. Configure the router to respond to ARP requests on the interface.

```
[edit interfaces demux0 unit 10303]
user@host1# set proxy-arp
```

b. Configure the outer VLAN tag to identify the access node (S-VLAN) and the inner VLAN tag to identify the subscriber port on the access node (C-VLAN).

```
[edit interfaces demux0 unit 10303]
user@host1# set vlan tags outer 103 inner 3
```

c. Specify that the VLAN demux interface runs on the underlying aggregated Ethernet interface.

```
[edit interfaces demux0 unit 10303]
user@host1# set demux-options underlying-interface ae0
```

d. Prevent multiple PPPoE sessions from being created for the same PPPoE subscriber on this VLAN demux interface.

```
[edit interfaces demux0 unit 10303]
user@host1# set family pppoe duplicate-protection
```

e. Configure the dynamic profile that is instantiated on the VLAN demux interface.

```
[edit interfaces demux0 unit 10303]
user@host1# set family pppoe dynamic-profile ancp-10303
```

f. Configure the recommended upstream and downstream traffic rates.

```
[edit interfaces demux0 unit 10303]
user@host1# set advisory-options upstream-rate 1m
user@host1# set advisory-options downstream-rate 16m
```

11. Configure the IPv4 protocol family and the address of the loopback interface.

 Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces lo0]
user@host1# set unit 0 family inet address 100.0.0.1/32
```

Results From configuration mode, confirm the static VLAN demux configuration by entering the **show interfaces** command.

```
[edit]
user@host# show interfaces
ge-1/0/1 {
   hierarchical-scheduler;
   gigether-options {
      802.3ad {
         ae0;
         primary;
      }
   }
}
ge-1/0/2 {
   hierarchical-scheduler;
   gigether-options {
      802.3ad {
         ae0;
         backup;
      }
   }
}
ae0 {
   hierarchical-scheduler;
   stacked-vlan-tagging;
   aggregated-ether-options {
      link-protection;
   }
}
demux0 {
   unit 10301 {
      proxy-arp;
      vlan-tags outer 103 inner 1;
      demux-options {
         underlying-interface ae0;
      }
      family pppoe {
         duplicate-protection;
         dynamic-profile ancp-10301;
      }
      advisory-options {
         downstream-rate 16m;
         upstream-rate 1m;
      }
   }
   unit 10302 {
      proxy-arp;
      vlan-tags outer 103 inner 2;
      demux-options {
         underlying-interface ae0;
      }
```

```
        family pppoe {
            duplicate-protection;
            dynamic-profile ancp-10302;
        }
        advisory-options {
            downstream-rate 16m;
            upstream-rate 1m;
        }
    }
    unit 10303 {
        proxy-arp;
        vlan-tags outer 103 inner 3;
        demux-options {
            underlying-interface ae0;
        }
        family pppoe {
            duplicate-protection;
            dynamic-profile ancp-10303;
        }
        advisory-options {
            downstream-rate 16m;
            upstream-rate 1m;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 100.0.0.1/32;
            }
        }
    }
}
```

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

Configuring Class of Service

Step-by-Step Procedure

1. Configure the traffic control profile with the shaping rate and specify the overhead accounting mode to account for ATM cell encapsulation.

```
[edit class-of-service]
user@host1# set traffic-control-profiles tcp1 shaping-rate 16m
user@host1# set traffic-control-profiles tcp1 overhead-accounting cell-mode
```

2. Apply the traffic control profile to the interface sets.

```
[edit class-of-service]
user@host1# set interfaces interface-set 10301 output-traffic-control-profile tcp1
user@host1# set interfaces interface-set 10302 output-traffic-control-profile tcp1
user@host1# set interfaces interface-set 10303 output-traffic-control-profile tcp1
```

Results

From configuration mode, confirm the class of service configuration by entering the **show class-of-service** command.

```
[edit]
user@host# show class-of-service
```


Copyright © 2013, Juniper Networks, Inc.

```
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 10.0.0.1
   ```

2. Configure the ANCP agent to report adjusted downstream traffic rates to CoS.

   ```
   [edit protocols ancp]
   user@host1# set qos-adjust
   ```

3. Specify an overhead adjustment of the traffic on ADSL and ADSL2 lines to 90 percent and 95 percent, respectively, of the net data rate.

   ```
   [edit protocols ancp]
   user@host1# set qos-adjust-adsl 90
   user@host1# set qos-adjust-adsl2 95
   ```

4. Specify an interval of 5 seconds between adjacency messages sent to all ANCP neighbors.

   ```
   [edit protocols ancp]
   user@host1# set adjacency-timer 5
   ```

5. Associate the ACI with the interface sets for each C-VLAN.

   ```
   [edit protocols ancp]
   user@host1# set interfaces interface-set 10301 access-identifier
      "192.168.61.65/0.0.0.0 eth 1/1:7;"
   user@host1# set interfaces interface-set 10302 access-identifier
      "192.168.61.65/0.0.0.0 eth 1/2:7;"
   user@host1# set interfaces interface-set 10303 access-identifier
      "192.168.61.65/0.0.0.0 eth 1/3:7;"
   ```

6. Specify the underlying interface for the interface sets.

   ```
   [edit protocols ancp]
   user@host1# set interfaces interface-set 10301 underlying-interface demux0.10301
   ```

```
user@host1# set interfaces interface-set 10302 underlying-interface demux0.10302
user@host1# set interfaces interface-set 10303 underlying-interface demux0.10303
```

7. Configure the size of the ANCP trace log files.

```
[edit protocols ancp traceoptions]
user@host1# set file ancpd size 512m
```

8. Configure flags for tracing ANCP configuration and CoS operations.

```
[edit protocols ancp traceoptions]
user@host1# set flag config
user@host1# set flag cos
```

Results From configuration mode, confirm the ANCP agent configuration by entering the **show ancp** command.

```
[edit]
user@host# show ancp
traceoptions {
   file ancpd size 512m;
   flag config;
   flag cos;
}
qos-adjust;
adjacency-timer 5;
qos-adjust-adsl 90;
qos-adjust-adsl2 95;
interfaces {
   interface-set {
      10301 {
         access-identifier "192.168.61.65/0.0.0.0 eth 1/1:7;";
         underlying-interface demux0.10301;
      }
      10302 {
         access-identifier "192.168.61.65/0.0.0.0 eth 1/2:7;";
         underlying-interface demux0.10302;
      }
      10303 {
         access-identifier "192.168.61.65/0.0.0.0 eth 1/3:7;";
         underlying-interface demux0.10303;
      }
   }
}
neighbor 10.0.0.1;
```

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 172.28.32.191 secret
   "$9$MUeL7VgoGqmTwYmTz3tpWLx"
```

2. Specify that RADIUS is used to authenticate subscribers.

 Copyright © 2013, Juniper Networks, Inc.

```
[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 172.28.32.191
user@host1# set profile radius-profile radius accounting-server 172.28.32.191
```

4. Configure options for the RADIUS server: The format used to identify the accounting session and that Juniper Networks DSL VSAs are added to RADIUS request messages.

```
[edit access]
user@host1# set profile radius-profile radius options accounting-session-id-format
  decimal
user@host1# set profile radius-profile radius options juniper-dsl-attributes
```

5. Exclude DSL Forum VSA attributes from being included in RADIUS messages.

```
[edit access]
user@host1# set profile radius-profile radius attribute exclude dsl-forum-attributes
  access-request
user@host1# set profile radius-profile radius attribute exclude dsl-forum-attributes
  accounting-start
user@host1# set profile radius-profile radius attribute exclude dsl-forum-attributes
  accounting-stop
```

Results From configuration mode, confirm the RADIUS configuration by entering the **show access** command.

```
[edit]
user@host# show access
radius-server {
   172.28.32.191 secret "$9$MUeL7VgoGqmTwYmTz3tpWLx"; ## SECRET-DATA
}
profile radius-profile {
   radius {
      authentication-server 172.28.32.191;
      accounting-server 172.28.32.191;
      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 1366
- Verifying the Traffic Scheduling and Shaping Parameters for the Interface Set on page 1366
- Verifying the demux0 Interface Configuration on page 1367
- Verifying the pp0 Interface Configuration on page 1367
- Verifying the ANCP Agent Configuration on page 1368

Verifying the Aggregated Ethernet Interface Configuration

Purpose Verify that the interface values match your configuration, the link is up, and traffic is flowing.

Action From operational mode, enter the **show interfaces redundancy** command.

```
user@host> show interfaces redundancy
Interface  State           Last change  Primary      Secondary    Current status
ae0        On primary                   ge-1/0/1     ge-1/0/2     both up
```

From operational mode, enter the **show interfaces ae0** command.

```
user@host> show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 606
  Link-level type: Ethernet, MTU: 1522, Speed: 1Gbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:1f:12:b8:ef:c0, Hardware address: 00:1f:12:b8:ef: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.

 Copyright © 2013, Juniper Networks, Inc.

Action `user@host>` **show class-of-service**

Verifying the demux0 Interface Configuration

Purpose Verify that the VLAN demux interface displays the configured PPPoE family attributes and the member links in the aggregated Ethernet bundle.

Action From operational mode, enter the **show interfaces demux0** command for each VLAN.

```
user@host> show interfaces demux0.10301
Logical interface demux0.10301 (Index 76) (SNMP ifIndex 61160)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
    Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 199)
    Link:
      ge-1/0/1
      ge-1/0/2
    Input packets : 2
    Output packets: 18575
    Protocol pppoe
      Dynamic Profile: ancp-10301,
      Service Name Table: None,
      Max Sessions: 16000, Duplicate Protection: On,
      AC Name: pppoe-server-1
```

Alternatively, you can enter **show pppoe underlying-interfaces detail** to display the state and PPPoE family configuration for all configured underlying interfaces.

Meaning The output shows the name of the underlying interface, the member links of the aggregated bundle, and the PPPoE family configuration. The output shows packet counts when traffic is present on the logical interface.

Verifying the pp0 Interface Configuration

Purpose Verify that the interface values match your configuration.

Action From operational mode, enter the **show interfaces pp0** command.

```
user@host> show interfaces pp0.100
  Logical interface pp0.100 (Index 71) (SNMP ifIndex 710)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 1,
      Session AC name: pppoe-server-1, Remote MAC address: 00:90:1a:00:18:34,
      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: 45.63.24.1
```

Meaning This output shows information about the PPPoE logical interface created on the underlying VLAN demux interface. The output includes the PPPoE family and aggregated Ethernet redundant link information, and shows input and output traffic for the PPPoE interface.

Verifying the ANCP Agent Configuration

Purpose Verify that the ANCP values match your configuration and that traffic is flowing.

Action From operational mode, enter the **show ancp subscriber** command.

```
user@host> show ancp subscriber detail
Interface  State           Last change  Primary      Secondary    Current status
ae0        On primary                   ge-1/0/1     ge-1/0/2     both up
```

From operational mode, enter the **show ancp cos** command.

```
user@host> show ancp cos

Qos Adjust Flag:         TRUE
Keepalive Timer:         30 secs
Cos State:               WRITE_READY
Connect Time:            Mon Mar 19 15:03:01 2012
Session Time:            Mon Mar 19 15:03:13 2012
Routing Instance Time:   Mon Mar 19 15:03:14 2012
Keepalive Time:          Not Set
Rate Update Time:        Mon Mar 19 15:03:15 2012


Type      Name             Index         Pending Update      Last Update
iflset    10301              1             None                   64 Kbps
iflset    10302              2             None                   64 Kbps
iflset    10303             71             None                   64 Kbps
```

Meaning The **show ancp subscriber** output shows subscriber line information such as state and the various traffic rates collected by the ANCP agent—displayed for each subscriber as identified by the ACI. The **show ancp cos** output shows that the ANCP agent is configured to send adjusted rate data to CoS, that keepalives are configured for a 30-second interval, and that the interface sets 10301, 10302, and 10303 are configured and their traffic rates are updating

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring a Dynamic Profile for Client Access on page 679
- Subscriber Interfaces and Demultiplexing Overview on page 759
- ANCP Interactions with AAA on page 1317
- ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes on page 1327

 Copyright © 2013, Juniper Networks, Inc.

- Configuring ANCP on page 1332
- AAA Service Framework Overview on page 21

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 81

Dynamic IGMP Configuration Overview

- Dynamic IGMP Configuration Overview on page 1371

Dynamic IGMP Configuration Overview

The Internet Group Management Protocol (IGMP) is a host to router signaling protocol for IPv4 used to support IP multicasting. This protocol manages the membership of hosts and routers in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routers. Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have members.

Subscriber access supports the configuration of IGMP within the **dynamic profiles** hierarchy. By specifying IGMP statements within a dynamic profile, you can dynamically apply IGMP configuration when a subscriber connects to an interface using a particular access technology (DHCP), enabling the subscriber to access a carrier (multicast) network.

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring a Dynamic Profile for Client Access on page 679
- For general information about configuring IGMP, see the *Multicast Protocols Feature Guide for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 82

Dynamic MLD Configuration Overview

- Dynamic MLD Configuration Overview on page 1373

Dynamic MLD Configuration Overview

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each router maintains a list of host multicast addresses that have listeners for each subnet, as well as a timer for each address. However, the router does not need to know the address of the listeners—just the address of the hosts. The router provides addresses to the multicast routing protocol it uses; this ensures that multicast packets are delivered to all subnets where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) protocol.

Subscriber access supports the configuration of MLD within the **dynamic profiles** hierarchy for dynamically created interfaces. By specifying MLD statements within a dynamic profile, you can dynamically apply MLD configuration when a subscriber connects to an interface using a particular access technology (DHCP), enabling the subscriber to access a carrier (multicast) network.

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring a Dynamic Profile for Client Access on page 679
- For general information about configuring MLD, see the *Multicast Protocols Feature Guide for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 83

Dynamic Router Advertisement Overview

- Dynamic Router Advertisement Configuration Overview on page 1375

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 on page 638
- Configuring a Dynamic Profile for Client Access on page 679
- Configuring an Address-Assignment Pool for Router Advertisement on page 171

Copyright © 2013, Juniper Networks, Inc.

- For general information about configuring the Router Advertisement Protocol, see the *Junos OS Routing Protocols Library for Routing Devices*.

Copyright © 2013, Juniper Networks, Inc.

PART 18

Subscriber Access Examples

- Service Profile Examples on page 1379

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 84

Service Profile Examples

- Example: Configuring a Tiered Service Profile for Subscriber Access on page 1379

Example: Configuring a Tiered Service Profile for Subscriber Access

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

The profile contains three services:

- Gold—Subscribers that pay for this service are allocated 10M bandwidth for data, voice, and video services.
- Silver—Subscribers that pay for this service are allocated 5M bandwidth for data, voice, and video services.
- Bronze—Subscribers that pay for this service are allocated 1M bandwidth for the data service only.

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

To configure a profile for a tiered service:

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

```
interfaces {
  ge-2/0/0 {
    description subscribers;
    hierarchical-scheduler;
    stacked-vlan-tagging;
    unit 1 {
      vlan-tags outer 100 inner 100;
      family inet {
        unnumbered-address lo0.0 preferred-source-address 100.0.0.1;
      }
    }
    unit 2 {
      family inet {
        vlan-tags outer 101 inner 101;
        unnumbered-address lo0.0 preferred-source-address 100.0.0.1;
      }
```

Copyright © 2013, Juniper Networks, Inc.

```
        }
        unit 3 {
          vlan-tags outer 102 inner 102;
          family inet {
            unnumbered-address lo0.0 preferred-source-address 100.0.0.1;
          }
        }
      }
    }
```

2. Configure the static CoS parameters.

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

```
class-of-service {
  forwarding-classes {
    queue 0 data;
    queue 1 voice;
    queue 2 video;
  }
  scheduler-maps {
    bronze_service_smap {
      forwarding-class data scheduler data_sch;
    }
    silver_service_smap {
      forwarding-class data scheduler data_sch;
      forwarding-class voice scheduler silver_voice_sch;
      forwarding-class video scheduler silver_video_sch;
    }
    gold_service_smap {
      forwarding-class data scheduler data_sch;
      forwarding-class voice scheduler gold_voice_sch;
      forwarding-class video scheduler gold_video_sch;
    }
  }
  schedulers {
    data_sch {
      transmit-rate percent 20;
      buffer-size remainder;
      priority low;
    }
    silver_voice_sch {
      transmit-rate percent 30;
      buffer-size remainder;
      priority high;
    }
    silver_video_sch {
      transmit-rate percent 30;
      buffer-size remainder;
      priority medium;
    }
    gold_voice_sch {
      transmit-rate percent 40;
      buffer-size remainder;
      priority high;
    }
```

 Copyright © 2013, Juniper Networks, Inc.

```
        gold_video_sch {
          transmit-rate percent 40;
          buffer-size remainder;
          priority medium;
        }
      }
    }
```

3. Configure the dynamic profile for the service.

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

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

4. Configure access for the subscribers.

   The DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server. You use DHCP relay to obtain configuration parameters, including an IP address, for subscribers. In this example, one DHCP server, address 100.20.42.1, can be used by subscribers.

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

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

```
forwarding-options {
  dhcp-relay {
    server-group {
```

```
            service_provider_group {
                100.20.42.1;
            }
        }
        group subscriber_group {
            active-server-group service_provider_group;
            dynamic-profile subscriber_profile;
            interface ge-2/0/0.1;
            interface ge-2/0/0.2;
            interface ge-2/0/0.3;
        }
    }
}
```

Related Documentation

- For more information about configuring CoS for subscriber access, see CoS for Subscriber Access Overview on page 949

Copyright © 2013, Juniper Networks, Inc.

PART 19

Complete Configuration Statement Hierarchy and Summary of Statements for Subscriber Access

- Subscriber Access Statement Hierarchy on page 1385
- Subscriber Access Configuration Statements on page 1413

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 85

Subscriber Access Statement Hierarchy

- [edit access address-assignment] Hierarchy Level on page 1385
- [edit access domain] Hierarchy Level on page 1386
- [edit access gx-plus] Hierarchy Level on page 1387
- [edit access profile] Hierarchy Level on page 1387
- [edit access tunnel-profile] Hierarchy Level on page 1389
- [edit access tunnel-switch-profile] Hierarchy Level on page 1390
- [edit diameter] Hierarchy Level on page 1390
- [edit dynamic-profiles] Hierarchy Level on page 1391
- [edit forwarding-options dhcp-relay] Hierarchy Level on page 1398
- [edit interfaces radius-options] Hierarchy Level on page 1402
- [edit jsrc] Hierarchy Level on page 1402
- [edit protocols ancp] Hierarchy Level on page 1403
- [edit services captive-portal-content-delivery] Hierarchy Level on page 1403
- [edit services l2tp] Hierarchy Level on page 1404
- [edit services mobile-ip] Hierarchy Level on page 1405
- [edit services radius-flow-tap] Hierarchy Level on page 1406
- [edit system services dhcp-local-server] Hierarchy Level on page 1407
- [edit system services packet-triggered-subscribers] Hierarchy Level on page 1411
- [edit system services static-subscribers] Hierarchy Level on page 1411
- [edit system services subscriber-management] Hierarchy Level on page 1412

[edit access address-assignment] Hierarchy Level

```
access {
  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 {
      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;
            }
          }
        }
      }
      link pool-name;
    }
```

Related Documentation

- Address-Assignment Pools Overview on page 167
- Configuring Address-Assignment Pools on page 168

[edit access domain] Hierarchy Level

```
access {
  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;
      tunnel-switch-profile profile-name;
    }
    parse-direction (left-to-right | right-to-left);
  }
}
```

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Domain Mapping Overview on page 180
- Configuring a Domain Map on page 181

[edit access gx-plus] Hierarchy Level

```
access {
  gx-plus {
    global {
      include-ipv6;
      max-outstanding-requests number;
    }
    partition partition-name {
      diameter-instance instance-name;
      destination-host hostname;
      destination-realm realm;
    }
  }
}
```

Related Documentation

- Gx-Plus for Provisioning Subscribers Overview on page 541
- Configuring Gx-Plus on page 549

[edit access profile] Hierarchy Level

```
access {
  profile profile-name {
    accounting {
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
      coa-immediate-update;
      coa-no-override service-class-attribute;
      duplication;
      immediate-update;
      order [ accounting-method ];
      statistics (time | volume-time);
      update-interval minutes;
      wait-for-acct-on-ack;
    }
    authentication-order [ authentication-methods ];
    authorization-order jsrc;
    client client-name {
      ...
    }
    domain-name-server;
    domain-name-server-inet;
    domain-name-server-inet6;
    provisioning-order (gx-plus | jsrc);
    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;
    }
    client-accounting-algorithm (detail | round-robin);
    client-authentication-algorithm(detail | round-robin);
    coa-dynamic-variable-validation;
    ethernet-port-type-virtual;
    interface-description-format {
        exclude-adapter;
        exclude-sub-interface;
    }
    nas-identifier identifier-value;
    nas-port-extended-format {
        adapter-width width;
        ae-width width;
        port-width width;
        slot-width width;
        stacked-vlan-width width;
        vlan-width width;
        atm {
            adapter-width width;
            port-width width:
            slot-width width;
            vci-width width:
            vpi-width width;
        }
    }
    nas-port-id-delimiter delimiter-character;
    nas-port-id-format {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        nas-identifier;
    }
    nas-port-type {
        ethernet {
            port-type;
        }
    }
    revert-interval interval;
```

 Copyright © 2013, Juniper Networks, Inc.

```
            vlan-nas-port-stacked-format;
          }
        }
        radius-server server-address {
          accounting-port port-number;
          port port-number;
          retry attempts;
          routing-instance routing-instance-name;
          secret password;
          max-outstanding-requests value;
          source-address source-address;
          timeout seconds;
        }
        service {
          accounting-order (activation-protocol | radius);
        }
        session-options {
          client-idle-timeout minutes;
          client-session-timeout minutes;
        }
      }
    }
```

Related Documentation

- AAA Service Framework Overview on page 21

[edit access tunnel-profile] Hierarchy Level

```
access {
  tunnel-profile profile-name {
    tunnel tunnel-id {
      identification name;
      logical-system logical-system-name;
      max-sessions number;
      medium type;
      nas-port-method cisco-avp;
      preference number;
      remote-gateway {
        address server-ip-address;
        gateway-name server-name;
      }
      routing-instance routing-instance-name;
      secret password;
      source-gateway {
        address client-ip-address;
        gateway-name client-name;
      }
      type tunnel-type;
    }
  }
}
```

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

[edit access tunnel-switch-profile] Hierarchy Level

```
access {
  tunnel-switch-profile profile-name {
    avp {
      bearer-type action;
      calling-number action;
      cisco-nas-port-info action;
    }
    tunnel-profile profile-name;
  }
}
```

Related Documentation

- Configuring L2TP Tunnel Switching on page 432

[edit diameter] Hierarchy Level

```
diameter {
  network-element element-name {
    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 {
    host hostname;
    realm realm-name;
  }
  peer peer-name {
    address ip-address;
    connect-actively {
      port port-number;
      transport transport-name;
    }
    logical-system logical-system-name <routing-instance routing-instance-name >;
    routing-instance routing-instance-name;
  }
  transport transport-name {
    address;
    logical-system logical-system-name <routing-instance routing-instance-name>;
    routing-instance routing-instance-name;
  }
}
```

Related Documentation

- Diameter Base Protocol Overview on page 455

Copyright © 2013, Juniper Networks, Inc.

- Configuring Diameter on page 471

[edit dynamic-profiles] Hierarchy Level

```
dynamic-profiles {
  profile-name {
    class-of-service {
      interfaces {
        interface-name {
          unit logical-unit-number {
            classifiers {
              type (classifier-name | default);
            }
            output-traffic-control-profile (profile-name |
              $junos-cos-traffic-control-profile);
            rewrite-rules {
              dscp (rewrite-name | default);
              dscp-ipv6 (rewrite-name | default);
              ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
              inet-precedence (rewrite-name | default);
              }
            }
          }
        }
      }
      scheduler-maps {
        map-name {
          forwarding-class class-name scheduler scheduler-name;
        }
      }
      schedulers {
        (scheduler-name) {
          buffer-size (percent percentage | remainder | temporal microseconds |
            $junos-cos-scheduler-bs);
          drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
            protocol (any | non-tcp | tcp) drop-profile (profile-name | predefined-variable);
          excess-priority (low | high | $junos-cos-scheduler-excess-priority);
          excess-rate (percent percentage  | percent $junos-cos-scheduler-excess-rate);
          overhead-accounting (shaping-mode) <bytes (byte-value>;
          priority (priority-level | $junos-cos-scheduler-priority);
          shaping-rate (rate | predefined-variable);
          transmit-rate (rate | percent percentage | remainder | percent percentage
            $junos-cos-scheduler-tx) <exact | rate-limit>;
        }
      }
      traffic-control-profiles profile-name {
        delay-buffer-rate (percent percentage | rate);
        excess-rate (percent percentage | proportion value | percent
          $junos-cos-excess-rate);
        guaranteed-rate (percent percentage | rate);
        overhead-accounting (shaping-mode) <bytes (byte-value>;
        scheduler-map map-name;
        shaping-rate (percent percentage | rate | predefined-variable);
      }
```

```
}
firewall {
  family family {
    fast-update-filter filter-name {
      interface-specific;
    match-order [match-order];
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      only-at-create;
    filter filter-name {
    interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
  policer policer-name {
    filter-specific;
    if-exceeding {
      (bandwidth-limit bps | bandwidth-percent percentage);
      burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    then {
      policer-action;
    }
  }
  hierarchical-policer policer-name {
    aggregate {
      if-exceeding {
        bandwidth-limit-limit bps;
        burst-size-limit bytes;
      }
      then {
        policer-action;
      }
    }
    premium {
      if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      then {
        policer-action;
```


Copyright © 2013, Juniper Networks, Inc.

```
            }
        }
    }
    three-color-policer policer-name {
        action {
            loss-priority high then discard;
        }
        logical-interface-policer;
        single-rate {
            (color-aware | color-blind);
            committed-burst-size bytes;
            committed-information-rate bps;
            excess-burst-size bytes;
        }
        two-rate {
            (color-aware | color-blind);
            committed-burst-size bytes;
            committed-information-rate bps;
            peak-burst-size bytes;
            peak-information-rate bps;
            }
        }
    }
}
policy-options {
    prefix-listname {
        ip-addresses;
    }
}
interfaces {
    interface-name {
        unit logical-unit-number {
            family family {
                access-concentrator name;
                address address;
                duplicate-protection;
                dynamic-profile profile-name;
                filter {
                    adf {
                        counter;
                        input-precedence precedence;
                        not-mandatory;
                        output-precedence precedence;
                        rule rule-value;
                    }
                    input filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                }
                max-sessions number;
                max-sessions-vsa-ignore;
```

```
            rpf-check {
              fail-filter filter-name;
              mode loose;
            }
            service {
              input {
                service-set service-set-name {
                  service-filter filter-name;
                }
                post-service-filter filter-name;
              }
              output {
                service-set service-set-name {
                  service-filter filter-name;
                }
              }
            }
            service-name-table table-name;
            short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
              maximum-seconds>;
            unnumbered-address interface-name <preferred-source-address address>;
          }
          ppp-options {
            chap;
            pap;
          }
          vlan-id number;
        }
        vlan-tagging;
      }
      interface-set interface-set-name {
        interface interface-name {
          unit logical-unit-number;
        }
      }
      demux0 {
        unit logical-unit-number {
          demux-options {
            underlying-interface interface-name
          }
          demux-source {
            source-prefix;
          }
          family family {
            access-concentrator name;
            address address;
            duplicate-protection;
            dynamic-profile profile-name;
            filter {
              input filter-name;
              output filter-name;
            }
            mac-validate (loose | strict):
            max-sessions number;
            max-sessions-vsa-ignore;
            service-name-table table-name;
```


Copyright © 2013, Juniper Networks, Inc.

```
                short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
                  maximum-seconds>;
                unnumbered-address interface-name <preferred-source-address address>;
              }
            }
          }
          pp0 {
            unit logical-unit-number {
              keepalives interval seconds;
              no-keepalives;
              pppoe-options {
                underlying-interface interface-name;
                server;
              }
              ppp-options {
                authentication [ authentication-protocols ];
                chap {
                  challenge-length minimum minimum-length maximum maximum-length;
                }
                pap;
              }
              family inet {
                unnumbered-address interface-name;
                address address;
                service {
                  input {
                    service-set service-set-name {
                      service-filter filter-name;
                    }
                    post-service-filter filter-name;
                  }
                  output {
                    service-set service-set-name {
                      service-filter filter-name;
                    }
                  }
                }
                filter {
                  input filter-name {
                    precedence precedence;
                  }
                  output filter-name {
                    precedence precedence;
                  }
                }
              }
            }
          }
        }
        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 {
                interface interface-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 {
                interface interface-name {
                  current-hop-limit number;
                  default-lifetime seconds;
                  (managed-configuration | no-managed-configuration);
                  max-advertisement-interval seconds;
                  min-advertisement-interval seconds;
                  (other-stateful-configuration | no-other-stateful-configuration);
                  prefix prefix {
                    (autonomous | no-autonomous);
                    (on-link | no-on-link);
                    preferred-lifetime seconds;
                    valid-lifetime seconds;
                  }
                  reachable-time milliseconds;
                  retransmit-timer milliseconds;
                }
              }
            }
          }
        }
        routing-instances routing-instance-name {
          interface interface-name;
```


Copyright © 2013, Juniper Networks, Inc.

```
    routing-options {
      access {
        route prefix {
          next-hop next-hop;
          metric route-cost;
          preference route-distance;
          tag route-tag;
        }
      }
      access-internal {
        route subscriber-ip-address {
          qualified-next-hop underlying-interface {
            mac-address address;
          }
        }
      }
      multicast {
        interface interface-name {
          no-qos-adjust;
        }
      }
    }
    rib routing-table-name {
      access {
        route prefix {
          next-hop next-hop;
          metric route-cost;
          preference route-distance;
          tag route-tag;
        }
      }
      access-internal {
        route subscriber-ip-address {
          qualified-next-hop underlying-interface {
            mac-address address;
          }
        }
      }
    }
  }
  routing-options {
    access {
      route prefix {
        next-hop next-hop;
        metric route-cost;
        preference route-distance;
        tag route-tag;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
```

```
        multicast {
          interface interface-name {
            no-qos-adjust;
          }
        }
      }
      variables {
        variable-name {
          default-value default-value;
          equals expression;
          mandatory;
          radius {
            vendor-id id {
              attribute attribute-number;
              tag tag-number;
            }
            redirect-url
          }
          uid;
          uid-reference;
        }
      }
    }
```

Related Documentation

- Dynamic Profiles Overview on page 638
- CoS for Subscriber Access Overview on page 949
- Configuring a Basic Dynamic Profile on page 673
- Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 963
- *Two-Color Policer Configuration Overview*
- *Three-Color Policer Configuration Overview*
- *Hierarchical Policer Configuration Overview*
- *Guidelines for Applying Traffic Policers*

[edit forwarding-options dhcp-relay] Hierarchy Level

```
forwarding-options {
  dhcp-relay {
    active-server-group server-group-name;
    authentication {
      password password-string;
      username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        logical-system-name;
        mac-address;
        option-60;
```


Copyright © 2013, Juniper Networks, Inc.

```
        option-82 <circuit-id> <remote-id>;
        routing-instance-name;
        user-prefix user-prefix-string;
      }
    }
    dhcpv6 {
      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-name;
          logical-system-name;
          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;
      }
      group group-name {
        active-server-group server-group-name;
        authentication {
          ...
        }
        dynamic-profile profile-name {
          ...
        }
        interface interface-name {
          exclude;
          overrides {
            ...
          }
          relay-option {
            ...
          }
          service-profile dynamic-profile-name;
          trace;
          upto upto-interface-name;
        }
        route-suppression;
        service-profile dynamic-profile-name;
        }
        overrides {
          ...
        }
        relay-agent-interface-id {
          ...
        }
```

```
        relay-option {
            ...
        }
        server-group {
            server-group-name {
                server-ip-address;
            }
        }
    }
    route-suppression;
    service-profile dynamic-profile-name;
    overrides {
        allow-snooped-clients;
        interface-client-limit number;
        no-allow-snooped-clients;
        no-bind-on-request;
        send-release-on-delete;
    }
    relay-agent-interface-id {
        prefix prefix;
        use-interface-description (logical | device);
        use-option-82;
    }
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
}
duplicate-clients-on-interface;
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 {
    active-server-group server-group-name;
    authentication {
        ...
    }
    dynamic-profile profile-name {
        ...
    }
    interface interface-name {
        exclude;
        overrides {
            ...
        }
        service-profile dynamic-profile-name;
        trace;
        upto upto-interface-name;
    }
    overrides {
        ...
    }
```

 Copyright © 2013, Juniper Networks, Inc.

```
    relay-option {
        ...
    }
    relay-option-82 {
        ...
    }
    route-suppression;
    service-profile dynamic-profile-name;
}
overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match <option60-and-option82>;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
```

```
        route-suppression;
        service-profile dynamic-profile-name;
      }
    }
```

Related Documentation

- Extended DHCP Relay Agent Overview on page 276

[edit interfaces radius-options] Hierarchy Level

```
interfaces interface-name {
  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);
    }
  }
}
```

Related Documentation

- RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN Overview on page 63
- Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 66

[edit jsrc] Hierarchy Level

```
jsrc {
  partition partition-name {
    diameter-instance instance-name;
    destination-host hostname;
    destination-realm realm-name;
  }
}
```

Related Documentation

- Juniper Networks Session and Resource Control (SRC) and JSRC Overview on page 485
- Configuring JSRC on page 491

Copyright © 2013, Juniper Networks, Inc.

[edit protocols ancp] Hierarchy Level

```
protocols {
  ancp {
    adjacency-timer seconds;
    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-timer;
      ietf-mode;
      maximum-discovery-table-entries entry-number;
      pre-ietf-mode;
    }
    pre-ietf-mode;
    qos-adjust;
    qos-adjust-adsl adjustment-factor;
    qos-adjust-adsl2 adjustment-factor;
    qos-adjust-adsl2-plus adjustment-factor;
    qos-adjust-sds1 adjustment-factor;
    qos-adjust-vds1 adjustment-factor;
    qos-adjust-vds2 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;
    }
  }
}
```

Related Documentation

- ANCP Topology Discovery and Traffic Reporting Overview on page 1311
- Configuring ANCP on page 1332

[edit services captive-portal-content-delivery] Hierarchy Level

```
services {
  captive-portal-content-delivery {
    rule rule-name {
      match-direction (input | output | input-output);
      term term-name {
        from {
          application [ application-name];
```

```
            destination-address address <except>;
            destination-prefix-list list-name <except>;
          }
          then {
            action;
            <action-modifiers>;
          }
        }
      }
      rule-set rule-set-name{
        [ rule rule-names];
      }
    }
  }
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit services] Hierarchy Level*

[edit services l2tp] Hierarchy Level

NOTE: The **tunnel-group** *group-name* stanza is not supported for L2TP LAC. It applies only to L2TP LNS. Similarly, some of the options for the **traceoptions** statement apply only to L2TP LNS; for more information, see traceoptions.

```
services {
  l2tp {
    destination
      lockout-timeout seconds;
    }
    destruct-timeout seconds;
    disable-calling-number-avp;
    disable-failover-protocol;
    enable-snmp-tunnel-statistics;
    fail-over-within-preference;
    ip-reassembly;
    rx-connect-speed-when-equal;
    traceoptions {
      debug-level level;
      file filename <files number> <match regular-expression > <size maximum-file-size>
        <world-readable | no-world-readable>;
      filter {
        protocol name;
        user-name username;
      }
      flag flag;
      interfaces interface-name {
        debug-level severity;
        flag flag;
      }
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
```


Copyright © 2013, Juniper Networks, Inc.

```
        }
        tunnel {
          assignment-id-format (assignment-id |client-server-id);
          idle-timeout seconds;
          retransmission-count-established count;
          retransmission-count-not-established count;
          tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port);
        }
        tunnel-group group-name {
          aaa-access-profile profile-name;
          dynamic-profile;
          hello-interval seconds;
          hide-avps;
          l2tp-access-profile profile-name;
          local-gateway address address;
          maximum-send-window packets;
          ppp-access-profile profile-name;
          receive-window packets;
          retransmit-interval seconds;
          service-device-pool;
          service-interface interface-name;
          syslog {
            host hostname {
              facility-override facility-name;
              log-prefix prefix-value;
              services severity-level;
            }
          }
          tos-reflect;
          tunnel-switch-profile profile-name;
          tunnel-timeout seconds;
        }
        tunnel-switch-profile profile-name;
        tx-connect-speed-method method;
        weighted-load-balancing;
      }
    }
```

Related Documentation

- L2TP for Subscriber Access Overview on page 379
- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

[edit services mobile-ip] Hierarchy Level

```
services {
  mobile-ip {
    access-type {
      (generic | wimax);
    }
    authenticate {
      order (aaa | local);
    }
    dynamic-home-assignment {
```

```
            home-agent {
              nai (name@domain | @domain) {
                home-agent ip-address;
              }
            }
          }
          home-agent {
            enable-service interface-name;
            virtual-network {
              home-agent-address ip-address {
                registration-lifetime seconds;
                revocation-required;
                timestamp-tolerance seconds;
              }
            }
          }
          peer {
            (ip-address address | nai name@domain) {
              spi hexadecimal-value {
                algorithm (hmac-md5 | md5);
                entity-type (host | mobility-agent);
                key (hex | ascii) string;
                replay-method (none | timestamp seconds);
              }
            }
          }
          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

- Mobile IP Home Agent Elements and Behavior on page 557
- Configuring Mobile IP on page 569

[edit services radius-flow-tap] Hierarchy Level

```
services {
  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;
```

 Copyright © 2013, Juniper Networks, Inc.

```
                dscp dscp-value;
                protocol protocol;
                source-address address;
                source-port port-number;
              }
            }
          }
          inet6 {
            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;
              }
            }
          }
        }
        source-ipv4-address ipv4-address;
      )
    }
```

Related Documentation

- Subscriber Secure Policy Overview on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

[edit system services dhcp-local-server] Hierarchy Level

```
system {
  services {
    dhcp-local-server {
      authentication {
        password password-string;
        username-include {
          circuit-type;
          delimiter delimiter-character;
          domain-name domain-name-string;
          interface-name;
          logical-system-name;
          mac-address;
          option-60;
          option-82 <circuit-id> <remote-id>;
          routing-instance-name;
          user-prefix user-prefix-string;
        }
      }
      dhcpv6 {
        authentication {
          password password-string;
          username-include {
            circuit-type;
```

```
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        logical-system-name;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix-string;
      }
    }
    group group-name {
      authentication {
        password password-string;
        username-include {
          circuit-type;
          client-id;
          delimiter delimiter-character;
          domain-name domain-name-string;
          interface-name;
          logical-system-name;
          relay-agent-interface-id;
          relay-agent-remote-id;
          relay-agent-subscriber-id;
          routing-instance-name;
          user-prefix user-prefix-string;
        }
      }
      interface interface-name {
        exclude;
        overrides {
          interface-client-limit number;
          process-inform {
            pool pool-name;
          }
          rapid-commit;
        }
        service-profile dynamic-profile-name;
        trace;
        upto upto-interface-name;
      }
      overrides {
        delegated-pool;
        interface-client-limit number;
        process-inform {
          pool pool-name;
        }
        rapid-commit;
      }
      reconfigure {
        attempts attempt-count;
        clear-on-abort;
        strict;
        timeout timeout-value;
        token token-value;
```

 Copyright © 2013, Juniper Networks, Inc.

```
            trigger {
              radius-disconnect;
            }
          }
          route-suppression;
          service-profile dynamic-profile-name;
        }
        overrides {
          delegated-pool;
          interface-client-limit number;
          process-inform {
            pool pool-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-on-interface;
      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 {
          password password-string;
          username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            overrides;
            routing-instance-name;
            user-prefix user-prefix-string;
            }
          }
        }
        dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
          primary-profile-name>;
        interface interface-name {
          exclude;
          overrides {
```

```
                client-discover-match <option60-and-option82>;
                interface-client-limit number;
                process-inform {
                  pool pool-name;
                }
              }
              service-profile dynamic-profile-name;
              trace;
              upto upto-interface-name;
            }
            overrides {
              client-discover-match <option60-and-option82>;
              interface-client-limit number;
              process-inform {
                pool pool-name;
              }
            }
            reconfigure {
              attempts attempt-count;
              clear-on-abort;
              timeout timeout-value;
              token token-value;
              trigger {
                radius-disconnect;
              }
            }
            route-suppression;
            service-profile dynamic-profile-name;
          }
          overrides {
            client-discover-match <option60-and-option82>;
            interface-client-limit number;
            process-inform {
              pool pool-name;
            }
            }
          service-profile dynamic-profile-name;
          pool-match-order {
            external-authority;
            ip-address-first;
            option-82;
          }
          reconfigure {
            attempts attempt-count;
            clear-on-abort;
            timeout timeout-value;
            token token-value;
            trigger {
              radius-disconnect;
            }
          }
        }
      }
      route-suppression;
      service-profile dynamic-profile-name;
    }
```


Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Extended DHCP Local Server Overview on page 200

[edit system services packet-triggered-subscribers] Hierarchy Level

```
system {
  services {
    packet-triggered-subscribers {
      subscriber-packet-idle-timeout subscriber-packet-idle-timeout
      partition partition-name {
        destination-host hostname;
        destination-realm realm;
        diameter-instance instance-name;
      }
      traceoptions {
        file filename <files number> <match regular-expression> <size maximum-file-size>
          <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
      }
    }
  }
}
```

Related Documentation

- Configuring the PTSP Application on page 525

[edit system services static-subscribers] Hierarchy Level

```
system {
  services {
    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>;
          }
          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 495
- Configuring Subscribers over Static Interfaces on page 500

[edit system services subscriber-management] Hierarchy Level

```
system {
  services {
    subscriber-management {
      enforce-strict-scale-limit-license;
      gres-route-flush-delay;
      maintain-subscriber {
        interface-delete;
      }
      traceoptions {
        file <filename> <files number> <match regular-expression > <size
          maximum-file-size> <world-readable | no-world-readable>;
        flag flag;
      }
    }
  }
}
```

Related Documentation

- Subscriber Binding Retention During Interface Delete Events on page 232
- Configuring the Router to Strictly Enforce the Subscriber Scaling License on page 235
- Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 693

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 86

Subscriber Access Configuration Statements

aaa-access-profile (L2TP LNS)

Syntax aaa-access-profile *profile-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*],
[edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced in Junos OS Release 11.4.
Support at the **[edit access profile *profile-name* client *client-name* l2tp]** hierarchy level introduced in Junos OS Release 12.1.

Description Specify a AAA access profile that overrides the AAA access profile configured for the routing instance with the **access-profile** statement. You can configure a profile to specify the RADIUS server settings for a tunnel group or for a LAC client, or both. The AAA access profile configured for the client takes precedence over the AAA access profile configured for the tunnel group, which takes precedence over the access profile configured for the routing instance.

Options ***profile-name***—Name of the local access profile for the tunnel group or client.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429

Copyright © 2013, 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.

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 184

Copyright © 2013, Juniper Networks, Inc.

aaa-routing-instance (Domain Map)

Syntax aaa-routing-instance *routing-instance-name*;

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.

Description Configure a non-default routing instance in which the **authd** daemon sends AAA requests for the domain map.

Default Default routing instance for the subscriber.

Options *routing-instance-name*—Name of the routing instance.

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 184

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 167
- Configuring Address-Assignment Pool Usage Threshold Traps on page 172

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 167
- Configuring Address-Assignment Pool Usage Threshold Traps on page 172

 Copyright © 2013, Juniper Networks, Inc.

access (Dynamic Access Routes)

Syntax

```
access {
  route prefix {
    next-hop next-hop;
    metric route-cost;
    preference route-distance;
    tag route-tag;
  }
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name*],
[edit dynamic-profiles *profile-name* routing-options]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name*]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure access routes.

BEST PRACTICE: We recommend that you always include the **access-internal** stanza in the dynamic-profile when the **access** stanza is present for framed-route support.

Options

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 690

access-concentrator

Syntax access-concentrator *name*;

Hierarchy Level
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-options],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information
Statement introduced before Junos OS Release 7.4.
Support at the **[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]** and **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]** hierarchy levels introduced in Junos OS Release 10.1.
Support at the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.

Description
(J Series Services Routers with Point-to-Point Protocol over Ethernet (PPPoE) interfaces) Configure the name of the access concentrator. If you configure a specific access concentrator name on the client and the same access concentrator name server is available, then a PPPoE session is established. If there is a mismatch between the access concentrator names of the client and the server, the PPPoE session gets closed.

If you do not configure the access concentrator name, the PPPoE session starts using any available server in the network.

(Intelligent Queuing 2 (IQ2) PICs on M120 and M320 routers; MPCs on MX Series routers) Configure an alternative access concentrator name in the AC-NAME tag in a PPPoE control packet for use with a dynamic PPPoE subscriber interface. If you do not configure the access concentrator name, the AC-NAME tag contains the system name.

NOTE: **The [edit ... family pppoe] hierarchies are supported only on MX Series routers with MPCs.**

Options *name*—Name of the access concentrator.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *Identifying the Access Concentrator*
- Configuring the PPPoE Family for an Underlying Interface on page 831
- Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles on page 899
- *Junos OS Interfaces and Routing Configuration Guide*

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 ANCP on page 1332
- Associating an Access Node with Subscribers for ANCP Operations on page 1334

access-internal (Dynamic Access-Internal Routes)

Syntax

```
access-internal {
  route subscriber-ip-address {
    qualified-next-hop underlying-interface {
      mac-address address;
    }
  }
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name*],
[edit dynamic-profiles routing-options]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name*]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure access-internal routes. Access-internal routes are optional, but are used instead of access routes if the next-hop address is not specified in the Framed-Route Attribute [22] for IPv4 or the Framed-IPv6-Route attribute [99] for IPv6.

BEST PRACTICE: We recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed-route support.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 692
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 692

Copyright © 2013, Juniper Networks, Inc.

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 47
- Configuring RADIUS Server Parameters for Subscriber Access on page 36

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 182

access-profile (Routing Instances)

Syntax access-profile *profile-name*;

Hierarchy Level [edit],
[edit routing-instances *routing-instances-name*]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Specify the access profile for use by the master routing instance.

Options ***profile-name***—Name of the 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 Access Components for the DHCP Layer 3 Wholesale Network Solution*
- *Configuring Access Components for the PPPoE Wholesale Network Solution*
- *Configuring Address Server Elements for the Broadband Subscriber Management Solution*

 Copyright © 2013, 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 500
- Specifying the Static Subscriber Global Access Profile on page 501
- Specifying the Static Subscriber Group Access Profile on page 505

Copyright © 2013, 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. (Not supported in Junos OS Release 13.1.)

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 28
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24

 Copyright © 2013, Juniper Networks, Inc.

access-type

Syntax

```
access-type {
   (generic | wimax);
}
```

Hierarchy Level

[edit services mobile-ip],
[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* services mobile-ip],
[edit routing-instances *routing-instance-name* services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.5.

Description

Configure the access type for Mobile IP.

The remaining statements are explained separately.

Default

The generic access type is used by default.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Access Type for Mobile IP on page 573

accounting (Access Profile)

Syntax

```
accounting {
  accounting-stop-on-access-deny;
  accounting-stop-on-failure;
  address-change-immediate-update;
  coa-immediate-update;
  coa-no-override service-class-attribute;
  duplication;
  duplication-vrf {
    access-profile-name profile-name;
    vrf-name vrf-name;
  }
  immediate-update;
  order [ accounting-method ];
  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.

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 24
- Configuring Per-Subscriber Session Accounting on page 29
- Understanding RADIUS Accounting Duplicate Reporting on page 28

Copyright © 2013, Juniper Networks, Inc.

accounting (Dynamic IGMP Interface)

Syntax (accounting | no-accounting);

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*],

Release Information Statement introduced in Junos OS Release 9.2.

Description Enable or disable the collection of IGMP join and leave event statistics for dynamically created IGMP interfaces.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about recording IGMP join and leave events, see "*Recording IGMP Join and Leave Events*" in the *Multicast Protocols Feature Guide for Routing Devices*

accounting (Dynamic MLD Interface)

Syntax (accounting | no-accounting);

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Enable or disable the collection of MLD join and leave event statistics for a dynamic interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Recording MLD Join and Leave Events*

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.

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 38

accounting-order (Service Accounting)

Syntax accounting-order (activation-protocol | radius);

Hierarchy Level [edit access profile *profile-name* service]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify which method is used for reporting subscriber service accounting.

Default **activation-protocol**

Options **activation-protocol**—Send service accounting reports by means of the application that activates services, such as JSRC.

radius—Send service accounting reports 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.

Related Documentation
- Configuring Service Accounting with JSRC on page 494
- Service Accounting with JSRC on page 486

 Copyright © 2013, Juniper Networks, Inc.

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.

Description Configure the port number on which to contact the accounting server.

Options ***port-number***—Port number on which to contact the accounting server. Most RADIUS servers use port number 1813 (as specified in RFC 2866).

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 23
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- *Configuring RADIUS Authentication for L2TP*

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 24

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 47
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24

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 24

 Copyright © 2013, 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 client access fails AAA but the AAA server grants 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 24

action

Syntax

```
action {
  loss-priority high then discard;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall three-color-policer *name*],
[edit firewall three-color-policer *name*],
[edit logical-systems *logical-system-name* firewall three-color-policer *name*]

Release Information

Statement introduced in Junos OS Release 8.2.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... three-color-policer]** hierarchy level introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Discard traffic on a logical interface using tricolor marking policing.

NOTE: **This statement is supported only on IQ2 interfaces.**

The remaining statement is explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Three-Color Policer Configuration Overview*
- *Basic Single-Rate Three-Color Policers*
- *Basic Two-Rate Three-Color Policers*
- *Two-Color and Three-Color Logical Interface Policers*
- *Two-Color and Three-Color Physical Interface Policers*
- *Two-Color and Three-Color Policers at Layer 2*
- loss-priority high then discard on page 1766

Copyright © 2013, 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 276
- Configuring Active Server Groups on page 328
- Group-Specific DHCP Relay Options on page 292
- **dhcp-relay on page 1548**

address

Syntax address (*ip-address* | *ipv6-address*);

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information Statement introduced in Junos OS Release 9.2.
Support at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family*]** hierarchy level introduced in Junos OS Release 10.1.

Description Configure the interface address.

Options ***ip-address***—IPv4 address of the interface.

ipv6-address—IPv6 address of the interface. When configuring an IPv6 address on a dynamically created interface, use the ***$junos-ipv6–address*** dynamic variable.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- "Configuring the Protocol Family," in *Junos OS Network Interfaces Library for Routing Devices*.
- *Junos OS Administration Library for Routing Devices*

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 471
- Configuring Diameter Peers on page 473

 Copyright © 2013, 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 471

address (Tunnel Profile Remote Gateway)

Syntax address *server-ip-address*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id* remote-gateway]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IP address of the remote gateway device at the L2TP tunnel endpoint, the LNS.

Options ***server-ip-address***—IP address of the remote gateway device.

Default: 0.0.0.0.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 403

address (Tunnel Profile Source Gateway)

Syntax address *client-ip-address*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id* source-gateway]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IP address of the source gateway device at the local L2TP tunnel endpoint, the LAC. This value overrides the default address for the logical system or routing instance.

Options ***client-ip-address***—IP address of the source gateway device.

Default: 0.0.0.0.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

 Copyright © 2013, 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 {
    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;
      }
    }
    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: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

pool-name—Name assigned to an address-assignment pool.

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 167
- Configuring Address-Assignment Pools on page 168

- Configuring an Address-Assignment Pool for L2TP LNS with Inline Services on page 424
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

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 Address-Change-Update message to the RADIUS accounting server. Any change to this setting takes effect for all new subscriber logins. Existing subscribers are not impacted by this change except when the AAA daemon restarts.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- *Saving IPv4 Addresses for Dual-Stack PPP Subscribers*

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 184

 Copyright © 2013, 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.

Description Prevents IPv6 prefixes from being used more than once when AAA is used to supply IPv6 prefixes for router advertisement.

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 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.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Duplicate Prefix Protection for Router Advertisement*
- *Duplicate Prefix Protection for NDRA*

Copyright © 2013, Juniper Networks, Inc.

adf (Dynamic Firewalls)

Syntax

```
adf {
   counter;
   input-precedence precedence;
   not-mandatory;
   output-precedence precedence;
   rule rule-value;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter]

Release Information

Statement introduced in Junos OS Release 10.4.
Option **not-mandatory** introduced in Junos OS Release 12.2.

Description

Configure an Ascend-Data-Filter that the dynamic profile applies to a subscriber session.

Options

counter—Enable a counter that increments each time the Ascend-Data-Filter rule is used. Typically used for testing purposes.

not-mandatory—Suppress router from reporting an error when the RADIUS reply message does not include the $junos-adf-rule-v4 or $junos-adf-rule-v6 variable that is configured for the Ascend-Data-Filter in the dynamic profile. In this circumstance, the Ascend-Data-Filter is not created.

precedence—Precedence value that sets the order in which dynamic service filters are applied on the interface. The lower the precedence value, the higher the precedence that is given. The precedence setting is used in conjunction with the precedence settings of all dynamic service filters configured (not only Ascend-Data-Filters) on the same interface to establish the order. For example, the order also includes any configured **input *filter-name* precedence *precedence*** and **output *filter-name* precedence *precedence*** statements.

Range: 0 through 255

Default: 0

rule-value—Ascend-Data-Filter rule. You can specify either a Junos predefined variable that maps the Ascend-Data-Filter actions to Junos filter functionality or you can manually configure the Ascend-Data-Filter rule. The router supports two predefined variables depending on family type: **$junos-adf-rule-v4** for family **inet** and **$junos-adf-rule-v6** for family **inet6**.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Firewall Filters Overview on page 1132
- Classic Filters Overview on page 1133
- Basic Classic Filter Syntax on page 1135

Copyright © 2013, Juniper Networks, Inc.

- For general information about configuring firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*.

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 ANCP on page 1332
- Specifying the Interval Between ANCP Adjacency Messages on page 1335
- Configuring ANCP Neighbors on page 1333

adjustment-control-profiles

Syntax

```
adjustment-control-profiles {
  profile-name {
    application {
      ancp;
      radius-coa;
      pppoe-tags;
      }
    }
  }
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 13.2.

Description Configure the CoS adjustment control profile.

Options *profile-name*—Name of the adjustment control profile.

The remaining statements are explained separately.

Required Privilege Level interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- CoS Adjustment Control Profiles Overview on page 1084
- Configuring CoS Adjustment Control Profiles on page 1106
- Verifying the CoS Adjustment Control Profile Configuration on page 1107
- **application (Adjustment Control Profiles) on page 1459**
- **overhead-accounting (Dynamic Traffic Shaping) on page 1841**

 Copyright © 2013, Juniper Networks, Inc.

adjust-minimum (Dynamic Shaping and Scheduling)

Syntax adjust-minimum (*rate* | $junos-cos-adjust-minimum);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*],
[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *traffic-control-profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description For adjustments performed by the ANCP or multicast applications on EQ DPCs and MPC/MIC interfaces, specify the minimum shaping rate for an adjusted scheduler node. The node is associated with a traffic-control profile.

For adjustments performed by the multicast application on MPC/MIC interfaces, specify the minimum shaping rate for an adjusted queue. The queue is associated with a scheduler.

Options ***rate***—Minimum shaping rate for a node or a queue, in Mbps

$junos-cos-adjust-minimum—Junos OS predefined variable that is replaced with the minimum shaping rate for a node that is obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes on page 1097
- Configuring a Dynamic Shaping-Rate Adjustment for Queues on page 1099

adjust-percent (Dynamic Schedulers)

Syntax adjust-percent *percentage*;

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description For a MPC/MIC interface, determine the percentage of adjustment for the shaping rate of a queue.

Options ***percentage***—Percentage of the shaping rate to adjust.

Range: 0 through 100 percent

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Shaping-Rate Adjustment for Queues on page 1099

 Copyright © 2013, 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.

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 1339
- Configuring ANCP on page 1332
- Configuring the Method to Set the LAC Connection Speeds to the LNS on page 408

Copyright © 2013, Juniper Networks, Inc.

agent-circuit-identifier (Dynamic VLAN interface Sets)

Syntax

```
agent-circuit-identifier {
   dynamic-profile profile-name;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" auto-configure],
[edit interfaces *interface-name* unit *logical-unit-number* auto-configure]

Release Information

Statement introduced in Junos OS Release 12.2.

Description

Configure a static or dynamic underlying VLAN interface to enable dynamic VLAN subscriber interface creation based on agent circuit identifier information.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information on page 737
- Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information on page 739

Copyright © 2013, Juniper Networks, Inc.

aggregate (Hierarchical Policer)

Syntax

```
aggregate {
  if-exceeding {
    bandwidth-limit bandwidth;
    burst-size-limit burst;
  }
  then {
    discard;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall hierarchical-policer *name*],
[edit firewall hierarchical-policer]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... hierarchical-policer *name*]** hierarchy level introduced in Junos OS Release 11.4.

Description

On M40e, M120, and M320 edge routers with Flexible PIC Concentrator (FPC) input as FFPC and FPC output as SFPC, and on MX Series, T320, T640, and T1600 edge routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure an aggregate hierarchical policer.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Hierarchical Policer Configuration Overview*
- *Hierarchical Policers*
- bandwidth-limit (Hierarchical Policer) on page 1475
- burst-size-limit (Hierarchical Policer) on page 1484
- hierarchical-policer on page 1675
- if-exceeding (Hierarchical Policer) on page 1692
- premium on page 1894

Copyright © 2013, Juniper Networks, Inc.

aggregate-clients (DHCP Local Server)

Syntax

aggregate-clients (merge | replace);

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dynamic-profile *profile-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* dynamic-profile *profile-name*],
[edit logical-systems *logical-system-name* system services dhcp-local-server dynamic-profile *profile-name*],
[edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* dynamic-profile *profile-name*],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dynamic-profile *profile-name*],
[edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* dynamic-profile *profile-name*],
[edit system services dhcp-local-server dynamic-profile *profile-name*],
[edit system services dhcp-local-server 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.

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.

Not supported for IP demux subscriber interfaces.

Options

merge—Aggregate multiple clients 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

- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

Copyright © 2013, 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 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.

Description

Specify that the router merge (chain) client attributes such as firewall filters and CoS attributes or replace them when multiple client sessions exist on the same underlying VLAN. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

Not supported for IP demux subscriber interfaces.

Options

merge—Aggregate multiple client attributes for the same subscriber (logical interface)

replace—Replace the entire logical interface whenever a new client logs in to the network using the same VLAN logical interface

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- dhcp-relay on page 1548
- *dhcp-relay (EX Series Switches only)*
- *Understanding the Extended DHCP Relay Agent for EX Series Switches*
- *Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)*
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236
- Group-Specific DHCP Relay Options on page 292

Copyright © 2013, 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 500
- Enabling Multiple Subscribers on a VLAN Logical Interface for All Static Subscribers on page 502

Copyright © 2013, Juniper Networks, Inc.

algorithm

Syntax algorithm (hmac-md5 | md5);

Hierarchy Level
[edit logical-systems *logical-system-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* services mobile-ip peer nai *user@domain* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain* spi *hexadecimal-value*],
[edit routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain* spi *hexadecimal-value*],
[edit services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit services mobile-ip peer nai *user@domain* spi *hexadecimal-value*]

Release Information
Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Configure the algorithm used for authenticating Mobile IP messages.

Default hmac-md5

Options
hmac-md5—Specifies algorithm hmac-md5

md5—Specifies algorithm md5

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

 Copyright © 2013, Juniper Networks, Inc.

allow-snooped-clients

Syntax allow-snooped-clients;

Hierarchy Level
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name* overrides],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay group *group-name* interface *interface-name* overrides],
[edit forwarding-options dhcp-relay group *group-name* overrides],
[edit forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay ...],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay ...]

Release Information
Statement introduced in Junos OS Release 10.2.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 12.1.

Description
Explicitly enable DHCP snooping support on the router.

Use the statement at the **[edit ... dhcpv6]** hierarchy levels to explicitly enable snooping support on the router for DHCPv6 relay agent.

NOTE: In Junos OS Release 10.0 and earlier, DHCP snooping is *enabled* by default. In Release 10.1 and later, DHCP snooping is *disabled* by default.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Extended DHCP Relay Agent Overview on page 276
- Overriding the Default DHCP Relay Configuration Settings on page 293
- DHCP Snooping Support on page 299

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 276
- dhcp-relay on page 1548

 Copyright © 2013, 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 276

anchor-point (Pseudowire Subscriber Interfaces)

Syntax anchor-point *lt-device*;

Hierarchy Level [edit interfaces ps0]

Release Information Statement introduced in Junos OS Release 13.1.

Description Specify the logical tunnel (lt) interface that identifies the Packet Forwarding Engine that processes the pseudowire termination.

Options ***lt-device***—An lt device in the format **lt-*fpc/pic/port***

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Configuring a Pseudowire Subscriber Logical Interface Device on page 935

 Copyright © 2013, Juniper Networks, Inc.

ancp

Syntax

```
ancp {
  adjacency-timer seconds;
  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-timer;
    ietf-mode;
    maximum-discovery-table-entries entry-number;
    pre-ietf-mode;
  }
  pre-ietf-mode;
  qos-adjust;
  qos-adjust-adsl adjustment-factor;
  qos-adjust-adsl2 adjustment-factor;
  qos-adjust-adsl2-plus adjustment-factor;
  qos-adjust-sds1 adjustment-factor;
  qos-adjust-vds1 adjustment-factor;
  qos-adjust-vds2 adjustment-factor;
  traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
```

Hierarchy Level

[edit protocols]

Release Information

Statement introduced in Junos OS Release 9.4.

Description

Configure Junos OS ANCP features.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring ANCP on page 1332

ancp (Adjustment Control Profiles)

Syntax

```
ancp {
   priority priority;
   algorithm algorithm;
}
```

Hierarchy Level

[edit class-of-service adjustment-control-profiles *profile-name* application]

Release Information

Statement introduced in Junos OS Release 13.1.

Description

Configure the shaping rate adjustment controls for the ANCP application.

Options

priority—Priority of the ANCP application in the adjustment control profile.

Range: 1 through 10; 1 being the highest priority.

Default: 1

algorithm—Rate adjustment algorithm used by the ANCP application.

Values:

- adjust-never—Do not perform rate adjustments.
- adjust-always—Adjust the shaping rate unconditionally.
- adjust-less—Adjust the shaping rate if it is less than the configured value.
- adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
- adjust-greater—Adjust the shaping rate if it is greater than the configured value.
- adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.

Default: adjust-always

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- CoS Adjustment Control Profiles Overview on page 1084
- Configuring CoS Adjustment Control Profiles on page 1106
- Verifying the CoS Adjustment Control Profile Configuration on page 1107
- **adjustment-control-profiles on page 1442**
- **application (Adjustment Control Profiles) on page 1459**

Copyright © 2013, Juniper Networks, Inc.

application (Adjustment Control Profiles)

Syntax

```
application {
  ancp;
  radius-coa;
  pppoe-tags;
}
```

Hierarchy Level [edit class-of-service adjustment-control-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure which applications in the adjustment control profile can make shaping rate adjustments.

The remaining statements are explained separately.

Required Privilege Level interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- CoS Adjustment Control Profiles Overview on page 1084
- Configuring CoS Adjustment Control Profiles on page 1106
- Verifying the CoS Adjustment Control Profile Configuration on page 1107
- **adjustment-control-profiles on page 1442**

application

Syntax application *application-name*;

Hierarchy Level [edit services application-identification rule *rule-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Identify the application for inclusion in a rule.

Options ***application-name***—Identifier for the application.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Redirecting HTTP Requests Overview on page 1225
- *Configuring APPID Rules*

apply-groups (Subscriber Secure Policy)

Syntax apply-groups *group-name*;

Hierarchy Level [edit services radius-flow-tap policy *policy-name* inet drop-policy *rule-name* from],
[edit services radius-flow-tap policy *policy-name* inet6 drop-policy *rule-name* from]

Release Information Statement introduced in Junos OS Release 12.3.

Description Specify groups from which to inherit configuration data for the radius-flow-tap policy.

Options ***group-name***— Name of the group that inherits the configuration data.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

apply-groups-except (Subscriber Secure Policy)

Syntax apply-groups-except *group-name*;

Hierarchy Level [edit services radius-flow-tap policy *policy-name* inet drop-policy *rule-name* from],
[edit services radius-flow-tap policy *policy-name* inet6 drop-policy *rule-name* from]

Release Information Statement introduced in Junos OS Release 12.3.

Description Specify groups from which to inherit configuration data for the radius-flow-tap policy.

Options ***group-name***— Name of the group that does not inherit the configuration data.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

 Copyright © 2013, Juniper Networks, Inc.

assignment-id-format (L2TP LAC)

Syntax assignment-id-format (assignment-id | client-server-id);

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 11.4.

Description Set the format for the name used for a tunnel, the tunnel assignment ID.

NOTE: **Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing no services l2tp tunnel assignment-id-format.**

Default assignment-Id

Options

assignment-Id—The tunnel name corresponds to RADIUS attribute Tunnel-Assignment-Id [82].

client-server-id—The tunnel name is a combination of RADIUS attributes Tunnel-Client-Auth-Id [90], Tunnel-Server-Auth-Id [91], and Tunnel-Assignment-Id [82].

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting the Format for the Tunnel Name on page 410

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 Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 246

Copyright © 2013, Juniper Networks, Inc.

attribute

Syntax attribute *attribute-number*;

Hierarchy Level [edit dynamic-profiles *profile-name* variables radius vendor-id]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure a RADIUS attribute as a variable in a dynamic profile.

Options ***attribute-number***—Number of the RADIUS attribute.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring User-Defined CoS Variables in a Dynamic Service Profile on page 1000

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.

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 58

authenticate

Syntax

```
authenticate {
  order (aaa | local);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip],
[edit routing-instances *routing-instances-name* services mobile-ip],
[edit services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the authentication method performed for Mobile IP.

The remaining statement is explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Access Type for Mobile IP on page 573

Copyright © 2013, 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-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.

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 214

authentication (DHCP Relay Agent)

Syntax

```
authentication {
  password password-string;
  username-include {
    circuit-type;
    client-id;
    delimiter delimiter-character;
    domain-name domain-name-string;
    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 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.

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.

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 1548
- Using External AAA Authentication Services with DHCP on page 214

Copyright © 2013, Juniper Networks, Inc.

authentication (Static and Dynamic PPP)

Syntax authentication [*authentication-protocols*];

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options], [edit interfaces pp0 unit *unit-number* ppp-options]

Release Information Statement introduced in Junos OS Release 12.2.

Description Specify the order in which the router tries to negotiate PPP authentication protocols when verifying that a PPP client can access the network. By default, the router tries to negotiate Challenge Handshake Authentication Protocol (CHAP) authentication first, and then tries Password Authentication Protocol (PAP) authentication if the attempt to negotiate CHAP authentication is unsuccessful.

You can specify one or both authentication protocols. If you specify both CHAP and PAP in either order, you must enclose the set of protocol names within square brackets ([ ]).

Options ***authentication-protocols***—One or both of the following PPP authentication protocols:

- **chap**—Challenge Handshake Authentication Protocol
- **pap**—Password Authentication Protocol

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Controlling the Negotiation Order of PPP Authentication Protocols on page 367

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.

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 500
- Configuring the Static Subscriber Global Authentication Password on page 503
- Configuring the Static Subscriber Group Authentication Password on page 506

Copyright © 2013, 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 introduced in Junos OS Release 11.2.

Description Set the order in which the Junos OS tries different authentication methods when verifying that a client can access the router or switch. For each login attempt, the software tries the authentication methods in order, from first to last.

Default password

Options *authentication-methods*

- **none**—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: For subscriber access management, you must always specify the radius method. Subscriber access management does not support the password option (the default), and authentication fails when no method is specified.

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 24
- *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 36

authorization-order

Syntax authorization-order jsrc;

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure AAA to use JSRC in an SRC environment to request authorization from the SAE when verifying that a DHCP subscriber can access the router. When you include this statement, AAA ignores any configured authentication order settings. This statement is ignored for non-DHCP subscribers.

Options jsrc—Use JSRC application to communicate with the SAE for subscriber authorization. JSRC is the only application that is currently available.

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 491
- Authorizing Subscribers with JSRC on page 493

 Copyright © 2013, Juniper Networks, Inc.

auto-configure

Syntax

```
auto-configure {
  vlan-ranges {
    access-profile profile-name;
    authentication {
      password password-string;
      username-include{
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-18;
        option-37;
        option-82 <circuit-id> <remote-id>;
        radius-realm radius-realm-string;
        user-prefix user-prefix-string;
      }
    }
    dynamic-profile profile-name {
      accept (any | dhcp-v4 |dhcp-v6| inet | inet6 | pppoe);
      ranges (any | low-tag)–(any | high-tag);
    }
    override;
  }
  stacked-vlan-ranges {
    access-profile profile-name;
    authentication {
      password password-string;
      username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-18;
        option-37;
        option-82 <circuit-id> <remote-id>;
        radius-realm radius-realm-string;
        user-prefix user-prefix-string;
      }
    }
    dynamic-profile profile-name {
      accept (any | dhcp-v4 |dhcp-v6| inet | inet6 | pppoe);
      ranges (any | low-tag–high-tag),(any | low-tag–high-tag);
    }
    override;
  }
  remove-when-no-subscribers;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Enable the configuration of dynamic, auto-sensed VLANs.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring VLAN Interfaces to Use Dynamic Profiles on page 720

auto-configure (Dynamic VLAN Interface Sets)

Syntax

```
auto-configure {
  agent-circuit-identifier {
    dynamic-profile profile-name;
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"],
[edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 12.2.

Description Enable the configuration of dynamic, auto-sensed VLAN subscriber interfaces on a static or dynamic underlying VLAN interface.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information on page 737
- Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information on page 739

 Copyright © 2013, 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*

avp (L2TP Tunnel Switching)

Syntax

```
avp {
  bearer-type;
  calling-number;
  cisco-nas-port-info;
}
```

Hierarchy Level [edit access tunnel-switch-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the action taken on L2TP AVPs that are negotiated when the first session is created; these AVPs are contained in the L2TP packets that are switched by the tunnel switch profile.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring L2TP Tunnel Switching on page 432

bandwidth (Inline Services)

Syntax bandwidth (1g | 10g);

Hierarchy Level [edit chassis fpc *slot-number* pic *number* inline-services]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the amount of bandwidth reserved on each Packet Forwarding Engine for tunnel traffic using inline services.

Options **1g**—Reserves 1 Gbps of bandwidth for tunnel traffic.

10g—Reserves 10 Gbps of bandwidth for tunnel traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling Inline Service Interfaces on page 427
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

 Copyright © 2013, Juniper Networks, Inc.

bandwidth-limit (Hierarchical Policer)

Syntax bandwidth-limit *bps*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall hierarchical-policer aggregate if-exceeding],
[editdynamic-profiles *profile-name* firewall hierarchical-policer premium if-exceeding],
[edit firewall hierarchical-policer aggregate if-exceeding],
[edit firewall hierarchical-policer premium if-exceeding]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... if-exceeding]** hierarchy level introduced in Junos OS Release 11.4.

Description For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure the maximum average bandwidth for premium or aggregate traffic in a hierarchical policer.

Options *bps*—You can specify the number of bits per second either as a decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 32,000 through 50,000,000,000

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Hierarchical Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Single Token Bucket Algorithm*
- *Determining Proper Burst Size for Traffic Policers*
- aggregate (Hierarchical Policer) on page 1447
- burst-size-limit (Hierarchical Policer) on page 1484
- premium (Hierarchical Policer) on page 1894

Copyright © 2013, Juniper Networks, Inc.

bandwidth-limit (Policer)

Syntax bandwidth-limit *bps*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall policer *policer-name* if-exceeding],
[edit firewall policer *policer-name* if-exceeding],
[edit logical-systems *logical-system-name* policer *policer-name* if-exceeding]

Release Information Statement introduced before Junos OS Release 7.4.
Support at the **[edit dynamic-profiles ... if-exceeding]** hierarchy level introduced in Junos OS Release 11.4.

Description For a single-rate two-color policer, configure the bandwidth limit as a number of bits per second. Single-rate two-color policing uses the single token bucket algorithm to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with **low** packet loss priority (PLP) and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

NOTE: **This statement specifies the bandwidth limit as an absolute number of bits per second. Alternatively, for single-rate two-color policers only, you can use the bandwidth-percent *percentage* statement to specify the bandwidth limit as a percentage of either the physical interface port speed or the configured logical interface shaping rate.**

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allows bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options *bps*—You can specify the number of bits per second either as a decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: (M Series, MX Series, and T Series routers) 8000 through 100,000,000,000

Default: None.

 Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Two-Color Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Single Token Bucket Algorithm*
- *Determining Proper Burst Size for Traffic Policers*
- bandwidth-percent on page 1478
- burst-size-limit (Policer) on page 1485

bandwidth-percent

Syntax bandwidth-percent *percentage*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall policer *policer-name* if-exceeding],
[edit firewall policer *policer-name* if-exceeding],
[edit logical-systems *logical-system-name* policer *policer-name* if-exceeding]

Release Information Statement introduced before Junos OS Release 7.4.
Support at the **[edit dynamic-profiles ... if-exceeding]** hierarchy level introduced in Junos OS Release 11.4.

Description For a single-rate two-color policer, configure the bandwidth limit as a percentage value. Single-rate two-color policing uses the *single token bucket algorithm* to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with **low** packet loss priority and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

NOTE: This statement specifies the bandwidth limit as a percentage of either the physical interface port speed or the configured logical interface shaping rate. Alternatively, you can use the **bandwidth-limit** ***bps*** statement to specify the bandwidth limit as an absolute number of bits per second.

The function of the bandwidth limit is extended by the burst size (configured using the **burst-size-limit** ***bytes*** statement) to allow bursts of traffic up to a limit based on the overall traffic load:

- When a single-rate two-color policer is applied to the input or output traffic at an interface, the initial capacity for traffic bursting is equal to the number of bytes specified by this statement.
- During periods of relatively low traffic (traffic that arrives at or departs from the interface at overall rates below the token arrival rate), unused tokens accumulate in the bucket, but only up to the configured token bucket depth.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower

 Copyright © 2013, Juniper Networks, Inc.

priority. You apply a hierarchical policer to ingress Layer 2 traffic to allows bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

percentage—Traffic rate as a percentage of either the physical interface media rate or the logical interface configured shaping rate. You can configure a shaping rate on a logical interface by using class-of-service statement.

NOTE: **You cannot rate-limit based on bandwidth percentage for aggregate, tunnel, and software interfaces. The bandwidth percentage policer cannot be used for forwarding table filters. Bandwidth percentage policers can only be used for interface-specific filters.**

Range: 0 through 100

Default: None.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Two-Color Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Single Token Bucket Algorithm*
- *Determining Proper Burst Size for Traffic Policers*
- *Bandwidth Policers*
- bandwidth-limit (Policer) on page 1476
- burst-size-limit (Policer) on page 1485

bearer-type (L2TP Tunnel Switching)

Syntax bearer-type *action*;

Hierarchy Level [edit access tunnel-switch-profile *profile-name* avp]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the action taken on the Bearer Type AVP (18) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.

Options ***action***—One of the following actions:

- **drop**—Drop the AVP.
- **regenerate**—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- **relay**—Forward the AVP transparently as is and send it in the switched packet.

Default: relay

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring L2TP Tunnel Switching on page 432

 Copyright © 2013, Juniper Networks, Inc.

bfd

Syntax

```
bfd {
  version (0 | 1 | automatic);
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  detection-time {
    threshold milliseconds;
  }
  session-mode (automatic | multihop | singlehop);
  holddown-interval milliseconds;
}
```

Hierarchy Level

[edit system services dhcp-local-server liveness-detection method],
[edit system services dhcp-local-server dhcpv6 liveness-detection method],
[edit forwarding-options dhcp-relay liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method],
[edit system services dhcp-local-server group *group-name* liveness-detection method],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure Bidirectional Forwarding Detection (BFD) as the liveness detection method.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

boot-file

Syntax boot-file *filename*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes]

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 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
- Configuring Address-Assignment Pools on page 168
- boot-server on page 1482

boot-server

Syntax boot-server (*address* | *hostname*);

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes]

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 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
- Configuring Address-Assignment Pools on page 168
- boot-file on page 1482

 Copyright © 2013, Juniper Networks, Inc.

buffer-size (Dynamic Scheduling)

Syntax

buffer-size (percent (*percentage* | $junos-cos-scheduler-bs) | remainder | temporal (*microseconds* | $junos-cos-scheduler-bs));

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]

Release Information

Statement introduced in Junos OS Release 9.3.
The **$junos-cos-scheduler-bs** predefined variable introduced in Junos OS Release 9.4.

Description

Specify buffer size.

Default

If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

Options

percent *percentage*—Buffer size as a percentage of total buffer.

remainder—Remaining buffer available.

temporal *microseconds*—Buffer size as a temporal value. The queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This maximum is computed by multiplying the logical interface speed by the configured temporal value.

Range: The ranges vary by platform as follows:

- For IQ PICs on M320 routers: 1 through 50,000 microseconds.
- For IQ PICs on other M Series routers: 1 through 100,000 microseconds.
- For other M Series routers: 1 through 200,000 microseconds.

$junos-scheduler-bs—Junos predefined variable that is replaced with the buffer size obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973
- **scheduler (Dynamic Scheduler Maps) on page 1968**

burst-size-limit (Hierarchical Policer)

Syntax burst-size-limit *bytes*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall hierarchical-policer aggregate if-exceeding],
[editdynamic-profiles *profile-name* firewall hierarchical-policer premium if-exceeding],
[edit firewall hierarchical-policer aggregate if-exceeding],
[edit firewall hierarchical-policer premium if-exceeding]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... if exceeding]** hierarchy level introduced in Junos OS Release 11.4.

Description For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs,T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure the burst-size limit for premium or aggregate traffic in a hierarchical policer.

Options *bytes*—Burst-size limit in bytes. The minimum recommended value is the maximum transmission unit (MTU) of the IP packets being policed. You can specify the value either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 2,147,450,880

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Hierarchical Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Single Token Bucket Algorithm*
- *Determining Proper Burst Size for Traffic Policers*
- *Hierarchical Policers*
- aggregate (Hierarchical Policer) on page 1447
- bandwidth-limit (Hierarchical Policer) on page 1475
- premium (Hierarchical Policer) on page 1894

 Copyright © 2013, Juniper Networks, Inc.

burst-size-limit (Policer)

Syntax burst-size-limit *bytes*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall policer *policer-name* if-exceeding],
[edit firewall policer *policer-name* if-exceeding],
[edit logical-systems *logical-system-name* policer *policer-name* if-exceeding]

Release Information Statement introduced before Junos OS Release 7.4.
Support at the **[edit dynamic-profiles ... if-exceeding]** hierarchy level introduced in Junos OS Release 11.4.

Description For a single-rate two-color policer, configure the burst size as a number of bytes. The burst size allows for short periods of traffic bursting (back-to-back traffic at average rates that exceed the configured bandwidth limit). Single-rate two-color policing uses the *single token bucket algorithm* to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with **low** packet loss priority and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

The burst size extends the function of the bandwidth limit (configured using either the **bandwidth-limit** ***bps*** statement or the **bandwidth-percent** ***percentage*** statement) to allow bursts of traffic up to a limit based on the overall traffic load:

- When a single-rate two-color policer is applied to the input or output traffic at an interface, the initial capacity for traffic bursting is equal to the number of bytes specified by this statement.
- During periods of relatively low traffic (traffic that arrives at or departs from the interface at overall rates below the token arrival rate), unused tokens accumulate in the bucket, but only up to the configured token bucket depth.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allows bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Copyright © 2013, Juniper Networks, Inc.

Table 123 on page 1486 summarizes the relationship between the **bandwidth-limit** and the token arrival rate. This information is useful in calculating the minimum **burst-size-limit**.

Table 123: Bandwidth Limits and Token Rates

Bandwidth Limit	Token Rate
0-333 Mbps	low (262 μs)
334-666 Mbps	high (8.2 μs)
667-1333 Mbps	low
1334 Mbps and above	high

The burst-size limit enforced is based on the burst-size limit you configure. For a rate-limited logical interface, the Packet Forwarding Engine calculates the optimum burst-size-limit values and then applies the value closest to the burst-size-limit value specified in the policer configuration.

On MX Series routers, the burst-size limit is not as freely configurable as it is on other platforms. Junos OS does not support an unlimited combination of policer bandwidth and burst-size limits on MX Series routers. For a single-rate two-color policer on an MX Series router, the minimum supported burst-size limit is equivalent to the amount of traffic allowed by the policer bandwidth limit in a time span of 1 millisecond. For example, for a policer configured with a **bandwidth-limit** value of 1 Gbps, the minimum supported value for **burst-size-limit** on an MX Series router is 125 KB. If you configure a value that is smaller than the minimum, Junos OS overrides the configuration and applies the actual minimum.

Options

bytes—Burst-size limit in bytes. The minimum recommended value is the maximum transmission unit (MTU) of the IP packets being policed. You can specify the value either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 100,000,000,000

Default: None

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *Two-Color Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Single Token Bucket Algorithm*
- *Determining Proper Burst Size for Traffic Policers*
- bandwidth-limit (Policer) on page 1476
- bandwidth-percent on page 1478

Copyright © 2013, Juniper Networks, Inc.

bytes (Dynamic Traffic Shaping)

Syntax bytes *bytes* | $junos-cos-byte-adjust;

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name* overhead-accounting],
[edit class-of-service traffic-control-profiles *profile-name* overhead-accounting]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the number of overhead bytes.

Options *bytes*—Byte adjustment value for the **cell-mode** or **frame-mode** shaping options. This can be the predefined variable **$junos-cos-byte-adjust**, which is the variable for byte adjustment that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

BEST PRACTICE: We recommend using the cell-mode cell-mode-bytes *cell-mode-bytes* option or the frame-mode frame-mode-bytes *frame-mode-bytes* option rather than the bytes option.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- CoS Adjustment Control Profiles Overview on page 1084
- Configuring CoS Adjustment Control Profiles on page 1106
- **adjustment-control-profiles on page 1442**
- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 1092
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072
- *egress-shaping-overhead*
- **cell-mode on page 1494**
- **frame-mode on page 1651**

 Copyright © 2013, Juniper Networks, Inc.

calling-number (L2TP Tunnel Switching)

Syntax calling-number *action*;

Hierarchy Level [edit access tunnel-switch-profile *profile-name* avp]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the action taken on the Calling Number AVP (22) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.

Options ***action***—One of the following actions:

- **drop**—Drop the AVP.
- **regenerate**—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- **relay**—Forward the AVP transparently as is and send it in the switched packet.

Default: relay

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring L2TP Tunnel Switching on page 432

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 56

 Copyright © 2013, Juniper Networks, Inc.

calling-station-id-format (Subscriber Management)

Syntax

```
calling-station-id-format {
   agent-circuit-id;
   agent-remote-id;
   interface-description;
   nas-identifier;
}
```

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 13.1.

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.

nas-identifier—Include the NAS-identifier (RADIUS IETF attribute 32), which specifies the name of the NAS that orginated the authentication or accounting request.

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 56

captive-portal-content-delivery

Syntax

```
captive-portal-content-delivery {
  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
      from {
        application [junos-http, junos-https, junos-httpproxy];
        destination-address address <except>;
        destination-prefix-list list-name <except>;
      }
      then {
        action;
        action-modifiers;
      }
    }
  }
  rule-set rule-set-name {
    [rule rule-names];
  }
}
```

Hierarchy Level

[edit services]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Configure the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

The remaining statements are explained separately.

Required Privilege Level

services—To view this statement in the configuration.
services–control—To add this statement to the configuration.

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

Copyright © 2013, Juniper Networks, Inc.

captive-portal-content-delivery-rule

Syntax captive-portal-content-delivery-rule *rule-name*;

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the HTTP rule for inclusion in a service set.

Options *rule-name*—Identifier for the rule.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Redirecting HTTP Requests Overview on page 1225

captive-portal-content-delivery-rule-set

Syntax captive-portal-content-delivery-rule-set *rule-set-name*;

Hierarchy Level [edit services (captive-portal-content-delivery)]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the HTTP rule set for inclusion in a service set.

Options *rule-set-name*—Identifier for the rule set.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Redirecting HTTP Requests Overview on page 1225

Copyright © 2013, Juniper Networks, Inc.

cell-mode (Dynamic Traffic Shaping)

Syntax

cell-mode (bytes *bytes* | $junos-cos-byte-adjust | cell-mode-bytes *cell-mode-bytes* |$junos-cos-byte-adjust-cell);

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name* overhead-accounting],
[edit class-of-service traffic-control-profiles *profile-name* overhead-accounting],

Release Information

Statement introduced in Junos OS Release 10.2.
Variable ***$junos-cos-byte-adjust-cell*** introduced in Junos OS Release 13.1.

Description

Configure the mode to shape downstream ATM traffic as cells.

Options

bytes—Byte adjustment value for the **cell-mode** or **frame-mode** shaping options.

$junos-cos-byte-adjust—Predefined variable for byte adjustment that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

cell-mode-bytes *cell-mode-bytes*—Shaping is based on the number of bytes in cells, and accounts for the ATM cell encapsulation and padding overhead. The resulting traffic stream conforms to the policing rates configured in downstream ATM switches, reducing the number of packet drops in the Ethernet network.

$junos-cos-byte-adjust-cell—Predefined variable for the cell mode shaping. This variable can not be used when the **overhead-accounting bytes *bytes*** option is configured.

BEST PRACTICE: We recommend using the **cell-mode-bytes *cell-mode-bytes*** option rather than the **bytes** option.

Range: –120 through 124 bytes

NOTE: If you specify a value for the **bytes *bytes*** option, you cannot specify a value for either the **cell-mode-bytes** option.

NOTE: Cell mode is supported only on logical interfaces and interface sets; it is not supported on physical interfaces (ifd or ifd-remaining).

Default: The default is frame-mode.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- CoS Adjustment Control Profiles Overview on page 1084
- Configuring CoS Adjustment Control Profiles on page 1106
- **adjustment-control-profiles on page 1442**
- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 1092
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072
- *egress-shaping-overhead*
- **bytes on page 1488**
- **frame-mode on page 1651**

challenge-length (Static and Dynamic PPP)

Syntax

challenge-length minimum *minimum-length* maximum *maximum-length*;

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options chap],
[edit interfaces pp0 unit *unit-number* ppp-options chap]

Release Information

Statement introduced in Junos OS Release 12.2.

Description

Modify the length of the Challenge Handshake Authentication Protocol (CHAP) challenge by specifying the minimum and maximum allowable length, in bytes.

BEST PRACTICE: **We recommend that you configure both the minimum length and the maximum length of the CHAP challenge to at least 16 bytes.**

Options

minimum-length—Minimum length, in bytes, of the CHAP challenge.

Range: 8 through 63

Default: 16

maximum-length—Maximum length, in bytes, of the CHAP challenge. The *maximum-length* must be equal to or greater than the *minimum-length*.

Range: 8 through 63

Default: 32

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Modifying the CHAP Challenge Length on page 369

chap

Syntax

```
chap {
   access-profile name;
   challenge-length minimum minimum-length maximum maximum-length;
   default-chap-secret name;
   local-name name;
   passive;
}
```

Hierarchy Level

[edit interfaces *interface-name* ppp-options],
[edit interfaces *interface-name* unit *logical-unit-number* ppp-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Allow each side of a link to challenge its peer, using a "secret" known only to the authenticator and that peer. The secret is not sent over the link.

By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

BEST PRACTICE: **On inline service (si) interfaces for L2TP, only the chap statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.**

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Junos OS Administration Library for Routing Devices*
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 419

Copyright © 2013, Juniper Networks, Inc.

chap (Dynamic PPP)

Syntax

```
chap {
   challenge-length minimum minimum-length maximum maximum-length;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]** hierarchy level introduced in Junos OS Release 12.2.

Description

Specify CHAP authentication in a PPP dynamic profile.

The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring Dynamic Authentication for PPP Subscribers on page 365
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 370
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 419

chap (L2TP)

Syntax

```
chap;
```

Hierarchy Level

[edit access group-profile *profile-name* ppp ppp-options]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

(MX Series routers only) Specify CHAP authentication for PPP subscribers in an L2TP LNS user group profile.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 418
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

Copyright © 2013, Juniper Networks, Inc.

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]

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 168

 Copyright © 2013, Juniper Networks, Inc.

circuit-id (DHCP Relay Agent)

Syntax

```
circuit-id {
  prefix prefix;
  use-interface-description (logical | device);
}
```

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.

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 or textual description, or both, instead of **circuit-id**.

The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual LANs (VLANs) or stacked VLANs (S-VLANs) is as follows:

(fe | ge)-*fpc/pic/port*

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*

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

- Enabling and Disabling Insertion of Option 82 Information on page 325
- Configuring Agent Circuit ID Information on page 325

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.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214

Copyright © 2013, Juniper Networks, Inc.

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 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.

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 214
- Creating Unique Usernames for DHCP Clients on page 239

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

Syntax cisco-nas-port-info *action*;

Hierarchy Level [edit access tunnel-switch-profile *profile-name* avp]

Release Information Statement introduced in Junos OS Release 13.2.

Description Define a tunnel profile for subscriber access.

Specify the action taken on the Cisco NAS Port Info AVP (100) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.

Options ***action***—One of the following actions:

- **drop**—Drop the AVP.
- **regenerate**—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- **relay**—Forward the AVP transparently as is and send it in the switched packet.

Default: relay

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring L2TP Tunnel Switching on page 432

 Copyright © 2013, Juniper Networks, Inc.

class-of-service (Dynamic Profiles)

Syntax class-of-service { ... }

Hierarchy Level [edit dynamic-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure Junos OS CoS features in a dynamic profile.

Default If you do not configure any CoS features, all packets are transmitted from output transmission queue 0.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 963
- Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access on page 965

classifiers (Dynamic CoS Application)

Syntax

```
classifiers {
  dscp (classifier-name | default);
  dscp-ipv6 (classifier-name | default);
  ieee-802.1 (classifier-name | default) vlan-tag (inner | outer)
  inet-precedence (classifier-name | default);
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Apply a CoS behavior aggregate classifier to a dynamic interface. You can apply a default classifier or one that is previously defined.

Options 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

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying a Classifier to a Subscriber Interface in a Dynamic Profile on page 981
- *classifiers (Definition)*

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 Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- Configuring Deletion of the Client When Dynamic Reconfiguration Fails on page 247

Copyright © 2013, Juniper Networks, Inc.

client

Syntax

```
client client-name {
  chap-secret chap-secret;
  group-profile profile-name;
  ike {
    allowed-proxy-pair {
      remote remote-proxy-address local local-proxy-address;
    }
    pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
    ike-policy policy-name;
    interface-id string-value;
  }
  l2tp {
    aaa-access-profile profile-name;
    interface-id interface-id;
    lcp-renegotiation;
    local-chap;
    maximum-sessions-per-tunnel number;
    multilink {
      drop-timeout milliseconds;
      fragment-threshold bytes;
    }
    ppp-authentication (chap | pap);
    ppp-profile profile-name;
    shared-secret shared-secret;
  }
  pap-password pap-password;
  ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-ip-address ip-address;
    framed-pool framed-pool;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
  }
  user-group-profile profile-name;
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the peer identity.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options ***client-name***—A peer identity. For L2TP clients, you can use a special name to configure a default client. This client enables the LNS to accept any LAC to establish the session. On M Series routers, use * for the default client configuration. On MX Series routers, use **default**.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- *Configuring the L2TP Client*
- *Configuring Access Profiles for L2TP or PPP Parameters*
- Configuring an L2TP Access Profile on the LNS on page 422

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.

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 36
- Configuring RADIUS Server Options for Subscriber Access on page 47

 Copyright © 2013, 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.

Description Configure the access method the router uses to access RADIUS authentication 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 36
- Configuring RADIUS Server Options for Subscriber Access on page 47

client-discover-match (DHCP Local Server)

Syntax

client-discover-match <option60-and-option82>;

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 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 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 group *group-name* overrides],
[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]

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure DHCP local server to use option 60 and option 82 information to uniquely identify DHCP subscribers or clients when primary subscriber or client identification fails. The statement always uses the **option60-and-option82** option. Specifying the option has no effect.

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 200
- Overriding Default DHCP Local Server Configuration Settings on page 220
- DHCP Auto Logout Overview on page 225

Copyright © 2013, Juniper Networks, Inc.

client-discover-match (DHCP Relay Agent)

Syntax

client-discover-match <option60-and-option82>;

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 9.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure DHCP relay to use option 60 and option 82 information to uniquely identify DHCP subscribers or clients when primary subscriber or client identification fails. The statement always uses the **option60-and-option82** option. Specifying the option has no effect.

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 276
- Overriding the Default DHCP Relay Configuration Settings on page 293
- DHCP Auto Logout Overview on page 225

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 239

Copyright © 2013, 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 214
- Creating Unique Usernames for DHCP Clients on page 239

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 (J Series, MX Series, and SRX Series devices) 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.

Default The timeout is not configured.

Options *minutes*—Number of minutes of idle time that elapse before the session is terminated.
Range: 10 through 1440 minutes

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 153
- Configuring Subscriber Session Options on page 154
- Removing Inactive Dynamic Subscriber VLANs on page 155

Copyright © 2013, 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 (J Series, MX Series, and SRX Series devices) Specify the amount of time after which user sessions are terminated, regardless of user activity (also known as a forced or hard authentication timeout).

Default The timeout is not configured.

Options *minutes*—Number of minutes after which user sessions are terminated.

Range: 1 through 527040 minutes

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 153
- Configuring Subscriber Session Options on page 154

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.

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 then 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 47
- RADIUS Server Options for Subscriber Access on page 41
- Configuring RADIUS Server Parameters for Subscriber Access on page 36

Copyright © 2013, Juniper Networks, Inc.

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 36
- Configuring Per-Subscriber Session Accounting on page 29

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 36
- Configuring Per-Subscriber Session Accounting on page 29

Copyright © 2013, Juniper Networks, Inc.

color-aware

Syntax color-aware;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall three-color-policer *name* single-rate],
[edit dynamic-profiles *profile-name* firewall three-color-policer *name* two-rate],
[edit firewall three-color-policer *policer-name* single-rate],
[edit firewall three-color-policer *policer-name* two-rate]

Release Information Statement introduced in Junos OS Release 7.4.
Support at the **[edit dynamic-profiles ... single-rate]** and **[edit dynamic-profiles ... two-rate]** hierarchy levels introduced in Junos OS Release 11.4.

Description For a three-color policer, configure the way preclassified packets are metered. In color-aware mode, the local router can assign a higher packet loss priority, but cannot assign a lower packet loss priority.

For example, suppose an upstream router assigned medium-high packet loss priority to a packet because the packet exceeded the committed information rate on the upstream router interface.

- If the local router applies color-aware policing to the packet, the router *cannot* change the packet loss priority to low, even if the packet conforms to the configured committed information route on the local router interface.
- If the local router applies color-blind policing to the packet, the router *can* change the packet loss priority to low if the packet conforms to the configured committed information route on the local router interface.

NOTE: **A color-aware policer cannot be applied to Layer 2 traffic.**

Default If you omit the **color-aware** statement, the default behavior is color-aware mode.

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Three-Color Policer Configuration Overview*
- *Color Modes for Three-Color Policers*
- color-blind on page 1515

 Copyright © 2013, Juniper Networks, Inc.

color-blind

Syntax color-blind;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall three-color-policer *name* single-rate],
[edit dynamic-profiles *profile-name* firewall three-color-policer *name* two-rate],
[edit firewall three-color-policer *policer-name* single-rate],
[edit firewall three-color-policer *policer-name* two-rate]

Release Information Statement introduced in Junos OS Release 7.4.
Support at the **[edit dynamic-profiles ... single-rate]** and **[edit dynamic-profiles ... two-rate]** hierarchy levels introduced in Junos OS Release 11.4.

Description For a three-color policer, configure the way preclassified packets are metered. In color-blind mode, the local router ignores the preclassification of packets and can assign a higher or lower packet loss priority.

For example, suppose an upstream router assigned medium-high packet loss priority to a packet because the packet exceeded the committed information rate on the upstream router interface.

- If the local router applies color-aware policing to the packet, the router *cannot* change the packet loss priority to low, even if the packet conforms to the configured committed information route on the local router interface.

  NOTE: **A color-aware policer cannot be applied to Layer 2 traffic.**

- If the local router applies color-blind policing to the packet, the router *can* change the packet loss priority to low if the packet conforms to the configured committed information route on the local router interface.

Default If you omit the **color-blind** statement, the default behavior is color-aware mode.

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Three-Color Policer Configuration Overview*
- *Color Modes for Three-Color Policers*
- color-aware on page 1514

committed-burst-size

Syntax committed-burst-size *bytes*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall three-color-policer *name* single-rate],
[edit dynamic-profiles *profile-name* firewall three-color-policer *name* two-rate],
[edit firewall three-color-policer *policer-name* single-rate],
[edit firewall three-color-policer *policer-name* two-rate]

Release Information Statement introduced in Junos OS Release 7.4.
Support at the **[edit dynamic-profiles ... single-rate]** and **[edit dynamic-profiles ... two-rate]** hierarchy levels introduced in Junos OS Release 11.4.

Description For a three-color policer, configure the committed burst size (CBS) as a number of bytes.

NOTE: When you include the **committed-burst-size** statement in the configuration, you must also include the **committed-information-rate** statement at the same hierarchy level.

In three-color policing, a committed information rate (CIR) defines the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions. A flow of traffic at an average rate that conforms to the CIR is categorized green.

During periods of average traffic rates below the CIR, any unused bandwidth capacity accumulates up to a maximum amount defined by the CBS. Short periods of bursting traffic (back-to-back traffic at averages rates that exceed the CIR) are also categorized as green provided that unused bandwidth capacity is available.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policers use a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red, based on the **excess-burst-size** statement included in the policer configuration.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits. Nonconforming traffic is categorized as yellow or red based on the **peak-information-rate** and **peak-burst-rate** statements included in the policer configuration.

Options *bytes*—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 100,000,000,000 bytes

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *Three-Color Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Dual Token Bucket Algorithms*
- *Determining Proper Burst Size for Traffic Policers*
- committed-information-rate on page 1518
- excess-burst-size on page 1619
- peak-burst-size on page 1857
- peak-information-rate on page 1859

Copyright © 2013, Juniper Networks, Inc.

committed-information-rate

Syntax committed-information-rate *bps*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall three-color-policer *name* single-rate],
[edit dynamic-profiles *profile-name* firewall three-color-policer *name* two-rate],
[edit firewall three-color-policer *policer-name* single-rate],
[edit firewall three-color-policer *policer-name* two-rate]

Release Information Statement introduced in Junos OS Release 7.4.
Support at the **[edit dynamic-profiles ... single-rate]** and **[edit dynamic-profiles ... two-rate]** hierarchy levels introduced in Junos OS Release 11.4.

Description For a three-color policer, configure the committed information rate as a number of bits per second. The committed information rate (CIR) is the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions.

NOTE: When you include the **committed-information-rate** statement in the configuration, you must also include the **committed-burst-size** statement at the same hierarchy level.

In three-color policing, a CIR defines the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions. A flow of traffic at an average rate that conforms to the CIR is categorized green.

During periods of average traffic rates below the CIR, any unused bandwidth capacity accumulates up to a maximum amount defined by the committed burst size (CBS). Short periods of bursting traffic (back-to-back traffic at averages rates that exceed the CIR) are also categorized as green provided that unused bandwidth capacity is available.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policers use a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red, based on the **excess-burst-size** statement included in the policer configuration.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits. Nonconforming traffic is categorized as yellow or red based on the **peak-information-rate** and **peak-burst-rate** statements included in the policer configuration.

Options *bps*—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 100,000,000,000 bps

 Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level
firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Three-Color Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Dual Token Bucket Algorithms*
- *Determining Proper Burst Size for Traffic Policers*
- committed-burst-size on page 1516
- excess-burst-size on page 1619
- peak-burst-size on page 1857
- peak-information-rate on page 1859

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.

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 471
- Configuring Diameter Peers on page 473

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*

 Copyright © 2013, Juniper Networks, Inc.

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.

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 161

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.

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 323

Copyright © 2013, 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*

default-value

Syntax default-value *default-value*;

Hierarchy Level [edit dynamic-profiles *profile-name* variables *variable-name*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure a default value for a user-defined variable in a dynamic profile. The values that the system uses for these variables are applied when the subscriber authenticates.

Options ***default-value***—Default value for the variable.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring User-Defined CoS Variables in a Dynamic Service Profile on page 1000

delay-buffer-rate (Dynamic Traffic Shaping)

Syntax delay-buffer-rate (percent *percentage* | *rate* | $junos-cos-delay-buffer-rate);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.2.
The **$junos-cos-delay-buffer-rate** variable introduced in Junos OS Release 9.4.

Description Base the delay-buffer calculation on a delay-buffer rate.

Default If you do not include this statement, the delay-buffer calculation is based on the guaranteed rate if one is configured, or the shaping rate if no guaranteed rate is configured.

Options *rate*—Delay-buffer rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1000 through 160,000,000,000 bps

$junos-cos-delay-buffer-rate—Junos predefined variable that is replaced with the delay-buffer rate obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Traffic Scheduling and Shaping for Subscriber Access on page 971
- output-traffic-control-profile on page 1839

 Copyright © 2013, Juniper Networks, Inc.

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 Pool for IPv6 Prefix Assignment on page 231
- Overriding Default DHCP Local Server Configuration Settings on page 220

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.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214

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 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.

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.

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 214
- Creating Unique Usernames for DHCP Clients on page 239

Copyright © 2013, Juniper Networks, Inc.

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 Name Delimiters on page 187
- Configuring Domain Name Usage for Domain Maps on page 187

demux0 (Dynamic Interface)

Syntax

```
demux0 {
  unit logical-unit-number {
    demux-options {
      underlying-interface interface-name
    }
    family family {
      access-concentrator name;
      address address;
      demux-source {
        source-prefix;
      }
      duplicate-protection;
      dynamic-profile profile-name;
      filter {
        input filter-name;
        output filter-name;
      }
      mac-validate (loose | strict):
      max-sessions number;
      max-sessions-vsa-ignore;
      rpf-check {
        fail-filter filter-name;
        mode loose;
      }
      service-name-table table-name
      short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
        maximum-seconds>;
      unnumbered-address interface-name <preferred-source-address address>;
    }
    filter {
      input filter-name;
      output filter-name;
    }
    vlan-id number;
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the logical demultiplexing (demux) interface in a dynamic profile.

Logical IP demux interfaces do not support IPv4 and IPv6 dual stack.

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 © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- For information about static IP demux interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

demux-options (Dynamic Interface)

Syntax

```
demux–options {
  underlying-interface interface-name
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces demux0 *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure logical demultiplexing (demux) interface options in a dynamic profile.

The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- For information about static IP demux interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

demux-source (Dynamic IP Demux Interface)

Syntax

```
demux-source {
   source-address;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure a logical demultiplexing (demux) source address for a subscriber in a dynamic profile.

Options

source-address—Either the specific source address you want to assign to the subscriber interface or the source address variable. For IPv4, specify **$junos-subscriber-ip-address**; for IPv6, specify **$junos-subscriber-ipv6-address**). The source address for the interface is dynamically supplied by DHCP when the subscriber accesses the router.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- For information about static IP demux interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

demux-source (Dynamic Underlying Interface)

Syntax demux-source *family*;

Hierarchy Level [edit dynamic-profiles interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the logical demultiplexing (demux) source family type on the IP demux underlying interface within a dynamic profile.

NOTE: **The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.**

Options *family*—Protocol family:

- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

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 471
- Configuring Diameter Network Elements on page 473

destination (L2TP Destination Lockout)

Syntax

```
destination {
  lockout-timeout seconds;
}
```

Hierarchy Level

[edit services l2tp]

Release Information

Statement introduced in Junos OS Release 13.2.

Description

Configure how long a destination is locked out from being considered when a new tunnel is created. Destinations are locked out when L2TP cannot connect to the destination during the tunnel selection process.

The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the L2TP Destination Lockout Timeout on page 414
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

destination-address

Syntax

destination-address *address* <except>;

Hierarchy Level

[edit services (captive-portal-content-delivery) captive-portal-content-delivery rule *rule-name* term *term-name* from]

Release Information

Statement introduced in Junos OS Release 10.4.
Option ***address*** enhanced to support IPv4 and IPv6 addresses in Junos OS Release 8.5.

Description

Specify the destination address for rule matching.

Options

address—Destination IPv4 or IPv6 address or prefix value.

except—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

Copyright © 2013, Juniper Networks, Inc.

destination-address (Subscriber Secure Policy)

Syntax destination-address *address*;

Hierarchy Level [edit services radius-flow-tap policy *policy-name* inet drop-policy *rule-name* from],
[edit services radius-flow-tap policy *policy-name* inet6 drop-policy *rule-name* from]

Release Information Statement introduced in Junos OS Release 12.3.

Description Specify destination IP address or prefix value for radius-flow-tap policy rule mapping.

Options ***address***— IPv4 or IPv6 address for the radius-flow-tap policy.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

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 491
- Configuring the JSRC Partition on page 492

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.

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 549
- Configuring the Gx-Plus Partition on page 550

destination-host (PTSP)

Syntax destination-host *hostname*;

Hierarchy Level [edit system services packet-triggered-subscribers partition *partition-name*]

Release Information Statement introduced in Junos OS Release 10.2.

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 the PTSP Partition on page 526

 Copyright © 2013, Juniper Networks, Inc.

destination-port (Subscriber Secure Policy)

Syntax destination-port *port-number*;

Hierarchy Level [edit services radius-flow-tap policy *policy-name* inet drop-policy *rule-name* from],
[edit services radius-flow-tap policy *policy-name* inet6 drop-policy *rule-name* from]

Release Information Statement introduced in Junos OS Release 12.3.

Description Specify the destination IP address for the radius-flow-tap policy.

Options ***port-number***— Number of the IPv4 or IPv6 destination port for the radius-flow-tap policy.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

destination-prefix-list

Syntax destination-prefix-list *list-name* <except>;

Hierarchy Level [edit services (captive-portal-content-delivery) captive-portal-content-delivery rule *rule-name* term *term-name* from]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the destination prefix list for rule matching. You configure the prefix list by including the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Options ***list-name***—Destination prefix list.

except—(Optional) Exclude the specified prefix list from rule matching.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Redirecting HTTP Requests Overview on page 1225
- *Routing Policy Feature Guide for Routing Devices*

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.

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 549
- Configuring the Gx-Plus Partition on page 550

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 491
- Configuring the JSRC Partition on page 492

 Copyright © 2013, Juniper Networks, Inc.

destination-realm (PTSP)

Syntax destination-realm *realm*

Hierarchy Level [edit system services packet-triggered-subscribers partition *partition-name*]

Release Information Statement introduced in Junos OS Release 10.2.

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 the PTSP Partition on page 526

destruct-timeout (L2TP)

Syntax destruct-timeout *seconds*;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 12.1.

Description Set how long the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed.

BEST PRACTICE: **Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing no services l2tp destruct-timeout.**

Options *seconds*—Length of the destruct timeout.

Range: 10 through 3600

Default: 300

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting the L2TP Destruct Timeout on page 413
- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

detection-time

Syntax

```
detection-time {
  threshold milliseconds;
}
```

Hierarchy Level

[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name*liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name*liveness-detection method bfd]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Enable failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

Copyright © 2013, Juniper Networks, Inc.

device-count (Pseudowire Subscriber Interfaces)

Syntax device-count *number*;

Hierarchy Level [edit chassis pseudowire-service]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the number of pseudowire logical devices available to the router.

Options *number*—Number of devices.
Range: 1 through 2048

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router on page 935

Copyright © 2013, 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;
  preferred-lifetime seconds
  valid-lifetime seconds
  t1-percentage percentage
  t2-percentage percentage
  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;
    }
  }
  router [ router-address ];
  server-identifier ip4-address;
  sip-server-address [ ipv6-address ];
  sip-server-domain-name domain-name;
  tftp-server address;
  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 address pools that can be used by different client applications.

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 167
- Configuring Address-Assignment Pools on page 168
- Configuring DHCP Client-Specific Attributes on page 172
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

Copyright © 2013, Juniper Networks, Inc.

dhcp-local-server

Syntax

```
dhcp-local-server {
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  dhcpv6 {
    authentication {
      ...
    }
    group group-name {
      authentication {
        ...
      }
      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);
              holddown-interval milliseconds;
            }
          }
        }
        overrides {
          interface-client-limit number;
          process-inform {
            pool pool-name;
          }
```

```
            rapid-commit;
        }
        service-profile dynamic-profile-name;
        trace;
        upto upto-interface-name;
    }
    overrides {
        delegated-pool;
        interface-client-limit number;
        process-inform {
            pool pool-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 {
    delegated-pool;
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
```


Copyright © 2013, Juniper Networks, Inc.

```
    route-suppression;
    service-profile dynamic-profile-name;
  }
  duplicate-clients-on-interface;
  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);
            holddown-interval milliseconds;
          }
        }
      }
      overrides {
        client-discover-match <option60-and-option82>;
        interface-client-limit number;
        process-inform {
          pool pool-name;
        }
      }
      service-profile dynamic-profile-name;
      trace;
      upto upto-interface-name;
    }
    overrides {
      client-discover-match <option60-and-option82>;
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
    }
    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 {
    client-discover-match <option60-and-option82>;
    interface-client-limit number;
    process-inform {
      pool pool-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;
    }
  }
  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]

Copyright © 2013, Juniper Networks, Inc.

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch and 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 so 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.

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 200
- DHCPv6 Local Server Overview on page 204
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

dhcp-relay

Syntax

```
dhcp-relay {
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  dhcpv6 {
    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-name;
        logical-system-name;
        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;
    }
    group group-name {
      active-server-group server-group-name;
      authentication {
        ...
      }
      dynamic-profile profile-name {
        ...
      }
      interface interface-name {
        exclude;
        liveness-detection {
          failure-action (clear-binding | clear-binding-if-interface-up | log-only);
          method {
```


Copyright © 2013, Juniper Networks, Inc.

```
                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 {
              ...
            }
            relay-option {
              ...
            }
            service-profile dynamic-profile-name;
            trace;
            upto upto-interface-name;
          }
          route-suppression:
          service-profile dynamic-profile-name;
          overrides {
            ...
          }
          relay-agent-interface-id {
            ...
          }
          relay-option {
            ...
          }
          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 {
    allow-snooped-clients;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
  }
  relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
  }
  server-group {
    server-group-name {
      server-ip-address;
    }
  }
duplicate-clients-on-interface;
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 {
  active-server-group server-group-name;
  authentication {
    ...
  }
  dynamic-profile 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;
```

 Copyright © 2013, Juniper Networks, Inc.

```
                }
                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;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode(automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
overrides {
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  client-discover-match <option60-and-option82>;
  disable-relay;
  interface-client-limit number;
  layer2-unicast-replies;
  no-allow-snooped-clients;
  no-bind-on-request;
```

```
    proxy-mode;
    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);
    }
  }
  server-group {
    server-group-name {
      server-ip-address;
    }
  }
  route-suppression:
  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.

Copyright © 2013, Juniper Networks, Inc.

Description

Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or 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 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 276
- DHCPv6 Relay Agent Overview on page 279
- DHCP Relay Proxy Overview on page 280
- Using External AAA Authentication Services with DHCP on page 214

Copyright © 2013, Juniper Networks, Inc.

dhcpv6 (DHCP Local Server)

Syntax

```
dhcpv6 {
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      logical-system-name;
      relay-agent-interface-id;
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  group group-name {
    authentication {
    ...
    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);
            holddown-interval milliseconds;
          }
        }
      }
      overrides {
        interface-client-limit number;
        process-inform {
          pool pool-name;
        }
        rapid-commit;
      }
      service-profile dynamic-profile-name;
      trace;
      upto upto-interface-name;
```


Copyright © 2013, Juniper Networks, Inc.

```
    }
    overrides {
      delegated-pool;
      interface-client-limit number;
      process-inform {
        pool pool-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 {
    delegated-pool;
    interface-client-limit number;
    process-inform {
      pool pool-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;
        }
    }
    route-suppression;
    service-profile dynamic-profile-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server],
[edit logical-systems *logical-system-name* system services dhcp-local-server],
[edit routing-instances *routing-instance-name* system services dhcp-local-server],
[edit system services dhcp-local-server]

Release Information

Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description

Configure DHCPv6 local server options on the router or switch and enable the router or switch to function as a server for the DHCP protocol for IP version 6 (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.

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 204

Copyright © 2013, Juniper Networks, Inc.

dhcpv6 (DHCP Relay Agent)

Syntax

```
dhcpv6 {
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      logical-system-name;
      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;
  }
  group group-name {
    active-server-group server-group-name;
    authentication {
      ...
    }
    dynamic-profile 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);
            holddown-interval milliseconds;
          }
        }
      }
```

```
        overrides {
            ...
        }
        service-profile dynamic-profile-name;
        trace;
        upto upto-interface-name;
    }
    }
    overrides {
        ...
    }
    relay-agent-interface-id {
        ...
    }
    relay-option {
        ...
    }
    route-suppression;
    service-profile dynamic-profile-name;
}
liveness-detection {
    ...
}
overrides {
    allow-snooped-clients;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
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;
    }
}
server-group {
    server-group-name {
        server-ip-address;
```

 Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  route-suppression;
  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.

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 1548
- DHCPv6 Relay Agent Overview on page 279
- Using External AAA Authentication Services with DHCP on page 214

dial-options

Syntax

```
dial-options {
   ipsec-interface-id name;
   l2tp-interface-id name;
   (shared | dedicated);
}
```

Hierarchy Level

[edit interfaces sp-*fpc/pic/port* unit *logical-unit-number*],
[edit interfaces si-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces sp-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces si-*fpc/pic/port* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
The **[edit ...si-...]** hierarchy levels introduced in Junos OS Release 11.4.

Description

Specify the options for configuring logical interfaces for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.

Options

dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.

ipsec-interface-id *name*—(M Series routers only) Interface identifier for group of dynamic peers. This identifier must be replicated at the **[edit access profile *name* client * ike]** hierarchy level.

l2tp-interface-id *name*—Interface identifier that must be replicated at the **[edit access profile *name*]** hierarchy level.

shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Identifier for Logical Interfaces that Provide L2TP Services*
- *Configuring Dynamic Endpoints for IPsec Tunnels*
- Configuring Options for the LNS Inline Services Logical Interface on page 428

Copyright © 2013, Juniper Networks, Inc.

dial-options (Dynamic Profiles)

Syntax

```
dial-options {
   ipsec-interface-id name;
   l2tp-interface-id name;
   (shared | dedicated);
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Specify the options for configuring logical interfaces in dynamic profiles for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.

Options

dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.

ipsec-interface-id ***name***—Interface identifier for group of dynamic peers. This identifier must be replicated at the **[edit access profile *name* client * ike]** hierarchy level. This options is not currently supported.

l2tp-interface-id ***name***—(MX Series routers only) L2TP interface identifier that must be replicated at the **[edit access profile *name*]** hierarchy level.

shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Dynamic LNS Sessions on page 431

diameter

Syntax

```
diameter {
  network-element element-name {
    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 {
    host hostname;
    realm realm-name;
  }
  peer peer-name {
    address ip-address;
    connect-actively {
      port port-number;
      transport transport-name;
    }
    logical-system logical-system-name <routing-instance routing-instance-name> ;
    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.

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 471

 Copyright © 2013, Juniper Networks, Inc.

diameter-instance (Gx-Plus)

Syntax diameter-instance *instance-name*;

Hierarchy Level [edit access gx-plus partition *partition-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the Diameter instance associated with the Gx-Plus 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 549
- Configuring the Gx-Plus Partition on page 550

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 491
- Configuring the JSRC Partition on page 492

Copyright © 2013, Juniper Networks, Inc.

diameter-instance (PTSP)

Syntax diameter-instance *instance-name*

Hierarchy Level [edit system services packet-triggered-subscribers partition *partition-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the Diameter instance associated with the PTSP 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 the PTSP Partition on page 526

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 © 2013, Juniper Networks, Inc.

disable (Dynamic IGMP)

Syntax "disable:$junos-igmp-enable";

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*],

Release Information Statement introduced in Junos OS Release 9.2.

Description Disable IGMP on the interface.

NOTE: Though the purpose of this statement is to disable IGMP on interfaces, under the **dynamic-profiles** hierarchy you can use this statement and an enable variable (**disable:$junos-igmp-enable**) to ensure that IGMP is not disabled by a AAA-based authentication and management method (RADIUS).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about disabling IGMP, see “Disabling IGMP” in the *Multicast Protocols Feature Guide for Routing Devices*

disable (Dynamic MLD)

Syntax disable;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Disable MLD on the dynamic interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Disabling MLD*

disable-calling-number-avp (L2TP LAC)

Syntax disable-calling-number-avp;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 10.4.

Description Prevent the LAC from sending L2TP Calling Number AVP 22 in incoming-call request (ICRQ) packets to the LNS. By default, the LAC in an L2TP network generates this AVP from the Calling-Station-Id and sends it to the LNS.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Preventing the LAC from Sending Calling Number AVP 22 to the LNS on page 407

disable-failover-protocol (L2TP LAC)

Syntax disable-failover-protocol;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure the LAC to use only the silent failover method when resynchronizing with the peer LNS in the event of LAC failover. This command prevents the default behavior, wherein the LAC first attempts to use the failover protocol and then falls back on the silent failover method. This configuration can be useful when routers that act as the LNS are configured for silent failover or incorrectly negotiate use of the failover protocol even though they do not support it.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Preventing the LAC From Negotiating L2TP Failover Protocol on page 409

 Copyright © 2013, 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 276

dns-server

Syntax dns-server *ipv6-address*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet6 dhcp-attributes]

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 167
- Configuring Address-Assignment Pools on page 168

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;
    tunnel-switch-profile profile-name;
  }
  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.

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 181

Copyright © 2013, Juniper Networks, Inc.

domain-name (Address-Assignment Pools)

Syntax domain-name *domain-name*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes]

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 168

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.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214

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 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.

Description

Specify the domain name that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

Options

domain-name-string—Domain name formatted string.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214
- Creating Unique Usernames for DHCP Clients on page 239

Copyright © 2013, 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 500
- Configuring the Static Subscriber Global Username on page 503
- Configuring the Static Subscriber Group Username on page 507

Copyright © 2013, 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.

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 lower in preference than one configured with the domain-name-server-inet 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 151
- DNS Name Server Address Overview on page 150

 Copyright © 2013, 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.

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.

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 151
- DNS Name Server Address Overview on page 150

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.

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 151
- DNS Name Server Address Overview on page 150

 Copyright © 2013, 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 1339
- Configuring ANCP on page 1332
- Configuring the Method to Set the LAC Connection Speeds to the LNS on page 408

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 323

Copyright © 2013, Juniper Networks, Inc.

drop-policy (Subscriber Secure Policy)

Syntax

```
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;
  }
}
```

Hierarchy Level

[edit services radius-flow-tap policy *policy-name* inet| inet6]

Release Information

Statement introduced in Junos OS Release 12.3.

Description

Specify the drop-policy that is applied to mirrored packets sent to a mediation device.

Options

rule-name—Define the term name.

The remaining statements are explained separately.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

drop-profile (Dynamic Schedulers)

Syntax drop-profile (*profile-name* | *predefined-variable*);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name* drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any | non-tcp | tcp)]

Release Information Statement introduced in Junos OS Release 9.3.
The **$junos-cos-scheduler-dropfile-low**, **$junos-cos-scheduler-dropfile-medium-low**, **$junos-cos-scheduler-dropfile-medium-high**, **$junos-cos-scheduler-dropfile-high**, and **$junos-cos-scheduler-dropfile-any** predefined variable introduced in Junos OS Release 9.4.

Description Within the drop-profile map, specify the name of the drop profile to use for random early detection (RED) for a specific packet-loss priority (PLP) level and protocol type. A drop profile maps a fill level (fullness of a queue) to a drop probability (probability that a packet will be dropped). When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.

You enable RED by applying a drop profile to a scheduler.

You configure drop profiles statically (at the **[edit class-of-service drop-profiles]** hierarchy level).

Options *profile-name*—Name of the drop profile.

predefined-variable—One of the following Junos predefined variable that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached:

- **$junos-cos-scheduler-dropfile-low**—Name of the drop profile for PLP level **low** and protocol **any**, specified for a scheduler configured in a dynamic profile for subscriber access.
- **$junos-cos-scheduler-dropfile-medium-low**—Name of the drop profile for PLP level **medium-low** and protocol **any**, specified for a scheduler configured in a dynamic profile for subscriber access.
- **$junos-cos-scheduler-dropfile-medium-high**—Name of the drop profile for PLP level **medium-high** and protocol **any**, specified for a scheduler configured in a dynamic profile for subscriber access.
- **$junos-cos-scheduler-dropfile-high**—Name of the drop profile for PLP level **high** and protocol **any**, specified for a scheduler configured in a dynamic profile for subscriber access.
- **$junos-cos-scheduler-dropfile-lny**—Name of the drop profile for PLP level **any** and protocol **any**, specified for a scheduler configured in a dynamic profile for subscriber access.

 Copyright © 2013, 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
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973
- **scheduler (Dynamic Scheduler Maps) on page 1968**
- For more information about configuring drop profiles and drop-profile maps, see the *Junos OS Class of Service Library for Routing Devices* .

drop-profile-map (Dynamic Schedulers)

Syntax
drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any | non-tcp | tcp) drop-profile (*profile-name* | *predefined-variable*);

Hierarchy Level
[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]

Release Information
Statement introduced in Junos OS Release 9.3.

Description
Define loss priority value for drop profile.

The 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
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973
- **scheduler (Dynamic Scheduler Maps) on page 1968**

dscp (Dynamic Classifiers)

Syntax dscp (*classifier-name* | default);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers]

Release Information Statement introduced in Junos OS Release 10.1.

Description For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) classifier to a subscriber interface in a dynamic profile.

Options *classifier-name*—Name of a **classifier** mapping configured at the **[edit class-of-service classifier dscp]** hierarchy level.

default—The default mapping.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying a Classifier to a Subscriber Interface in a Dynamic Profile on page 981
- *classifiers (Definition)*

 Copyright © 2013, Juniper Networks, Inc.

dscp (Dynamic Rewrite Rules)

Syntax dscp (*rewrite-name* | default);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]

Release Information Statement introduced in Junos OS Release 10.1.

Description For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) rewrite rule to a subscriber interface in a dynamic profile.

Options ***rewrite-name***—Name of a **rewrite-rules** mapping configured at the **[edit class-of-service rewrite-rules dscp]** hierarchy level.

default—The default mapping.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile on page 980
- *rewrite-rules*

dscp (Subscriber Secure Policy)

Syntax dscp *value*;

Hierarchy Level [edit services radius-flow-tap policy *policy-name* inet drop-policy *rule-name* from],
[edit services radius-flow-tap policy *policy-name* inet6 drop-policy *rule-name* from]

Release Information Statement introduced in Junos OS Release 12.3.

Description Specify the DSCP value for the radius-flow-tap policy.

Options ***dscp-value***— IPv4 or IPv6 dscp value for the radius-flow-tap policy.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

dscp-ipv6 (Dynamic Classifiers)

Syntax dscp-ipv6 (*classifier-name* | default);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers]

Release Information Statement introduced before Junos OS Release 10.1.

Description For IPv6 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) classifier to a subscriber interface in a dynamic profile.

Options *classifier-name*—Name of a **classifier** mapping configured at the **[edit class-of-service classifier ieee-802.1]** hierarchy level.

default—The default mapping.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying a Classifier to a Subscriber Interface in a Dynamic Profile on page 981
- *classifiers (Definition)*

dscp-ipv6 (Dynamic Rewrite Rules)

Syntax dscp-ipv6 (*rewrite-name* | default);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]

Release Information Statement introduced before Junos OS Release10.1.

Description For IPv6 traffic, apply a DSCP rewrite rule to a subscriber interface in a dynamic profile.

Options *rewrite-name*—Name of a **rewrite-rules** mapping configured at the **[edit class-of-service rewrite-rules dscp-ipv6]** hierarchy level.

default—The default mapping.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- *rewrite-rules*

 Copyright © 2013, Juniper Networks, Inc.

duplicate-clients-on-interface (DHCP Local Server)

Syntax duplicate-clients-on-interface;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server],
[edit logical-systems *logical-system-name* system services dhcp-local-server],
[edit routing-instances *routing-instance-name* system services dhcp-local-server],
[edit system services dhcp-local-server]

Release Information Statement introduced in Junos OS Release 10.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure DHCP local server to include the client subinterface when distinguishing between duplicate DHCP clients (clients with the same MAC address or client ID) in the same subnet. By default, DHCP distinguishes clients by subnet. This feature is supported on DHCPv4 only.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring DHCP Duplicate Client Support on page 214

duplicate-clients-on-interface (DHCP Relay Agent)

Syntax duplicate-clients-on-interface;

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 10.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure DHCP relay agent to include the client subinterface when distinguishing between duplicate DHCP clients (clients with the same MAC address or client ID) in the same subnet. By default, DHCP relay distinguishes clients by subnet. This feature is supported on DHCPv4 only.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring DHCP Duplicate Client Support on page 214
- Enabling and Disabling Insertion of Option 82 Information on page 325

duplicate-protection (Dynamic PPPoE)

Syntax duplicate-protection;

Hierarchy Level
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number*family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information
Statement introduced in Junos OS Release 10.1.
Support for the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.

Description Prevent the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same media access control (MAC) address is already active on that interface. Duplicate protection is disabled by default. Enabling duplicate protection has no effect on dynamic PPPoE logical interfaces that are already active.

NOTE: The **[edit ... family pppoe]** hierarchies are supported only on MX Series routers with MPCs.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces on page 905
- Configuring the PPPoE Family for an Underlying Interface on page 831
- Configuring Lockout of PPPoE Subscriber Sessions on page 912
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

 Copyright © 2013, 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 24
- Understanding RADIUS Accounting Duplicate Reporting on page 28

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.

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 28
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24

dynamic-class-of-service-options (Dynamic Traffic Shaping)

Syntax

```
dynamic-class-of-service-options {
   vendor-specific-tags access-loop-encapsulation;
   vendor-specific-tags actual-data-rate-downstream;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service]

Release Information

Statement introduced in Junos OS Release 12.1.

Description

Configure the shaping-rate and overhead-accounting class-of-service attributes based on access line parameters in PPPoE discovery packets on dynamic subscriber interfaces.

Options

vendor-specific-tags—Use Vendor-Specific Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101] to set the rate-shaping and overhead-accounting class-of-service attributes.

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

- Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 1074
- Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces on page 1093
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072

Copyright © 2013, Juniper Networks, Inc.

dynamic-home-assignment

Syntax

```
dynamic-home-assignment {
  home-agent {
    nai (name@domain.com | @domain.com) {
      home-agent ip-address;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip],
[edit routing-instances *routing-instances-name* services mobile-ip],
[edit services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the dynamic assignment rule for the home agent.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring Dynamic Home Assignment for the Mobile Node on page 572

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.

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 on page 236
- Configuring a Default Subscriber Service on page 1127

Copyright © 2013, 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 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.

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.

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 1548
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236
- Grouping Interfaces with Common DHCP Configurations on page 217
- Configuring a Default Subscriber Service on page 1127

Copyright © 2013, 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 183

 Copyright © 2013, Juniper Networks, Inc.

dynamic-profile (Dynamic PPPoE)

Syntax

dynamic-profile *profile-name*;

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information

Statement introduced in Junos OS Release 10.1.
Support for the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.

Description

Attach a PPPoE dynamic profile to an underlying Ethernet interface. This underlying interface is configured with either the **encapsulation ppp-over-ether** statement or the **family pppoe** statement; the two statements are mutually exclusive. When the router creates a dynamic PPPoE logical interface on the underlying interface, it uses the information in the dynamic profile to determine the properties of the dynamic PPPoE logical interface.

NOTE: The **[edit ... family pppoe]** hierarchies are supported only on MX Series routers with MPCs.

Options

profile-name—Name of a previously configured PPPoE dynamic profile, up to 64 characters in length, defined at the **[edit dynamic-profiles *profile-name* interfaces pp0]** hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces on page 905
- Configuring the PPPoE Family for an Underlying Interface on page 831
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

dynamic-profile (Dynamic VLAN Interface Sets)

Syntax dynamic-profile *profile-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" auto-configure agent-circuit-identifier],
[edit interfaces *interface-name* unit *logical-unit-number* auto-configure agent-circuit-identifier]

Release Information Statement introduced in Junos OS Release 12.2.

Description Attach a dynamic profile for an agent circuit identifier (ACI) interface set to a static or dynamic underlying VLAN interface.

Options
- *profile-name*—Name of the dynamic profile that defines the ACI interface set.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information on page 737
- Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information on page 739

dynamic-profile (L2TP)

Syntax dynamic-profile *profile-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Assign a dynamic profile to the tunnel group for dynamic LNS sessions.

Options ***profile-name***—Name of the dynamic profile for the tunnel group.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Dynamic Profile for Dynamic LNS Sessions on page 431

 Copyright © 2013, Juniper Networks, Inc.

dynamic-profile (PPP)

Syntax dynamic-profile *profile-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* ppp-options]

Release Information Statement introduced in Junos OS Release 9.5.
Support for MLPPP on LSQ interfaces introduced in Junos OS Release 10.2.

Description Specify the dynamic profile that is attached to the interface. On the MX Series routers, this statement is currently supported on PPPoE interfaces only. On the M120 and M320 routers, this statement is supported for MLPPP bundles only on LSQ interfaces on Adaptive Services PICs and Multiservices PICs.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring a Basic Dynamic Profile on page 673
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 370
- Attaching Dynamic Profiles to MLPPP Bundles on page 375
- For hardware requirements, see Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces on page 374

Copyright © 2013, Juniper Networks, Inc.

dynamic-profile (PPPoE Service Name Tables)

Syntax dynamic-profile *profile-name*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify a dynamic profile to instantiate a dynamic PPPoE interface. You can associate a dynamic profile with a named service entry, **empty** service entry, or **any** service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The dynamic profile associated with a service entry in a PPPoE service name table overrides the dynamic profile associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the **dynamic-profile** statement at the **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level, you cannot also include the **static-interface** statement at this level. The **dynamic-profile** and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options ***profile-name***—Name of the dynamic profile that the router uses to instantiate a dynamic PPPoE interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring PPPoE Service Name Tables*
- Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation on page 910

 Copyright © 2013, 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.

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 500
- Specifying the Static Subscriber Global Dynamic Profile on page 502
- Specifying the Static Subscriber Group Dynamic Profile on page 505

Copyright © 2013, Juniper Networks, Inc.

dynamic-profiles

Syntax

```
dynamic-profiles {
  profile-name {
    class-of-service {
      interfaces {
        interface-name ;
        }
        unit logical-unit-number {
          classifiers {
            type (classifier-name | default);
          }
          output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
          rewrite-rules {
            dscp (rewrite-name | default);
            dscp-ipv6 (rewrite-name | default);
            ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
            inet-precedence (rewrite-name | default);
            }
          }
        }
      }
      scheduler-maps {
        map-name {
          forwarding-class class-name scheduler scheduler-name;
        }
      }
      schedulers {
        (scheduler-name) {
          buffer-size (seconds | percent percentage | remainder | temporal microseconds);
          drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
            protocol (any | non-tcp | tcp) drop-profile profile-name;
          excess-priority (low | high | $junos-cos-scheduler-excess-priority);
          excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
          overhead-accounting (shaping-mode) <bytes (byte-value>;
          priority priority-level;
          shaping-rate (rate | predefined-variable);
          transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
        }
      }
      traffic-control-profiles profile-name {
        delay-buffer-rate (percent percentage | rate | $junos-cos-delay-buffer-rate);
        excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
        guaranteed-rate (percent percentage | rate | $junos-cos-guaranteed-rate);
        overhead-accounting (shaping-mode) <bytes (byte-value>;
        scheduler-map map-name;
        shaping-rate (rate | predefined-variable);
      }
    }
    firewall {
      family family {
        fast-update-filter filter-name {
          interface-specific;
          match-order [match-order];
```


Copyright © 2013, Juniper Networks, Inc.

```
    term term-name {
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
      only-at-create;
    }
  }
  firewall {
    family family {
      fast-update-filter filter-name {
        interface-specific;
      match-order [match-order];
        term term-name {
          from {
            match-conditions;
          }
          then {
            action;
            action-modifiers;
          }
        only-at-create;
      filter filter-name {
      interface-specific;
        term term-name {
          from {
            match-conditions;
          }
          then {
            action;
            action-modifiers;
          }
        }
    policer policer-name {
      filter-specific;
      if-exceeding {
        (bandwidth-limit bps | bandwidth-percent percentage);
        burst-size-limit bytes;
      }
      logical-bandwidth-policer;
      logical-interface-policer;
      physical-interface-policer;
      then {
        policer-action;
      }
    }
    hierarchical-policer policer-name {
      aggregate {
        if-exceeding {
          bandwidth-limit-limit bps;
          burst-size-limit bytes;
        }
        then {
```

```
                policer-action;
              }
            }
            premium {
              if-exceeding {
                bandwidth-limit bps;
                burst-size-limit bytes;
              }
              then {
                policer-action;
              }
            }
          }
          three-color-policer policer-name {
            action {
              loss-priority high then discard;
            }
            logical-interface-policer;
            single-rate {
              (color-aware | color-blind);
              committed-burst-size bytes;
              committed-information-rate bps;
              excess-burst-size bytes;
            }
            two-rate {
              (color-aware | color-blind);
              committed-burst-size bytes;
              committed-information-rate bps;
              peak-burst-size bytes;
              peak-information-rate bps;
              }
            }
          }
        }
        policy-options {
          prefix-listname {
            ip-addresses;
          }
        }
      }
    }
    interfaces interface-name {
      interface-set interface-set-name {
        interface interface-name {
          unit logical unit number {
            advisory-options {
              downstream-rate rate;
              upstream-rate rate;
            }
          }
        }
      }
      unit logical-unit-number {
        auto-configure {
          agent-circuit-identifier {
            dynamic-profile profile-name;
```

 Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid |
    atm-tcc-vc-mux | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux |
    atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc |
    ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet |
    frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type |
    frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end | multilink-ppp |
    ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc
    | vlan-tcc | vlan-vpls);
  family family {
    address address;
    filter {
      adf {
        counter;
        input-precedence precedence;
        not-mandatory;
        output-precedence precedence;
        rule rule-value;
      }
      input filter-name (
        precedence precedence;
      }
      output filter-name {
        precedence precedence;
      }
    }
    rpf-check {
      fail-filter filter-name;
      mode loose;
    }
    service {
      input {
        service-set service-set-name {
          service-filter filter-name;
        }
        post-service-filter filter-name;
      }
      input-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (push | swap);
        tag-protocol-id tpid;
        vlan-id number;
      }
      output {
        service-set service-set-name {
          service-filter filter-name;
        }
      }
      output-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (pop | swap);
        tag-protocol-id tpid;
        vlan-id number;
```

```
            }
          }
          unnumbered-address interface-name <preferred-source-address address>;
        }
        ppp-options {
          chap;
          pap;
        }
        vlan-id number;
        vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
      }
    }
    interfaces {
      demux0 {...}
    }
    interfaces {
      pp0 {...}
    }
    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;
```

 Copyright © 2013, Juniper Networks, Inc.

```
            }
        }
        router-advertisement {
            interface interface-name {
                current-hop-limit number;
                default-lifetime seconds;
                (managed-configuration | no-managed-configuration);
                max-advertisement-interval seconds;
                min-advertisement-interval seconds;
                (other-stateful-configuration | no-other-stateful-configuration);
                prefix prefix;
                reachable-time milliseconds;
                retransmit-timer milliseconds;
            }
        }
    }
}
routing-instances routing-instance-name {
    interface interface-name;
    routing-options {
        access {
            route prefix {
                next-hop next-hop;
                metric route-cost;
                preference route-distance;
                tag route-tag;
            }
        }
        access-internal {
            route subscriber-ip-address {
                qualified-next-hop underlying-interface {
                    mac-address address;
                }
            }
        }
        multicast {
            interface interface-name {
                no-qos-adjust;
            }
        }
    }
    rib routing-table-name {
        access {
            route prefix {
                next-hop next-hop;
                metric route-cost;
                preference route-distance;
                tag route-tag;
            }
        }
        access-internal {
            route subscriber-ip-address {
                qualified-next-hop underlying-interface {
                    mac-address address;
                }
            }
```

```
            }
          }
        }
        routing-options {
          access {
            route prefix {
              next-hop next-hop;
              metric route-cost;
              preference route-distance;
              tag route-tag;
            }
          }
          access-internal {
            route subscriber-ip-address {
              qualified-next-hop underlying-interface {
                mac-address address;
              }
            }
          }
          multicast {
            interface interface-name {
              no-qos-adjust;
            }
          }
        }
        variables {
          variable-name {
            default-value default-value;
            equals expression;
            mandatory;
            radius {
              vendor-id id {
                attribute attribute-number;
                tag tag-number;
              }
            }
            uid;
            uid-reference;
          }
        }
      }
    }
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.2.
Support at the **filter, policer, hierarchical-policer, three-color-policer, and policy options** hierarchy levels introduced in Junos OS Release 11.4.

Description Create dynamic profiles for use with DHCP or PPP client access.

Options ***profile-name***—Name of the dynamic profile; string of up to 80 alphanumeric characters.

The remaining statements are explained separately.

 Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring a Basic Dynamic Profile on page 673
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Dynamic Profiles Overview on page 638

effective-shaping-rate

Syntax
effective-shaping-rate;

Hierarchy Level
[edit chassis]

Release Information
Statement introduced in Junos OS Release 13.2.

Description
Specify that the Cos-Effective-Shaping-Rate VSA [26–177] included in RADIUS Acct-Start, Acct-Stop, and Interim-Acct messages reports the actual rate of the downstream traffic for a subscriber, in kilobits per second.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Reporting the Effective Shaping Rate for Subscribers on page 1108

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.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

enable-service

Syntax enable-service *interface-name*;

Hierarchy Level [edit logical-systems *logical-system-name* services mobile-ip home-agent],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent],
[edit services mobile-ip home-agent]

Release Information Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Define the list of interfaces on which the home agent service can be enabled. The system accepts registration requests only if it is on one of these interfaces. Include the statement once for each interface to be enabled.

Options ***interface-name***—Interface on which the home agent can be enabled.

Required Privilege Level view—To view this statement in the configuration.
view-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

 Copyright © 2013, Juniper Networks, Inc.

enable-snmp-tunnel-statistics (L2TP)

Syntax enable-snmp-tunnel-statistics;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 12.1R4 and supported in later 12.1Rx releases.
Statement supported in Junos OS Release 12.2R2 and later 12.2Rx releases. (Not supported in Junos OS Release 12.2R1.)
Statement supported in Junos OS Release 12.3 and later releases.

Description Enable collection of L2TP tunnel and global counters for SNMP statistics.

NOTE: The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.

Default Disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling Tunnel and Global Counters for SNMP Statistics Collection on page 416

encapsulation (Dynamic Interfaces)

Syntax encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type | frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls);

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Dynamic interface configuration of the logical link-layer encapsulation type.

Options **atm-ccc-cell-relay**—Use ATM cell-relay encapsulation.

atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on circuit cross-connect (CCC) circuits. When you use this encapsulation type, you can configure the **ccc** family only.

atm-cisco-nlpid—Use Cisco ATM network layer protocol ID (NLPID) encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink Point-to-Point Protocol (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a link services or voice services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-ppp-llc—For ATM2 IQ interfaces only, use PPP over AAL5 LLC encapsulation.

atm-ppp-vc-mux—For ATM2 IQ interfaces only, use PPP over ATM AAL5 multiplex encapsulation.

atm-snap—Use ATM subnetwork attachment point (SNAP) encapsulation.

atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

atm-vc-mux—Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

ether-over-atm-llc—For interfaces that carry IPv4 traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

 Copyright © 2013, Juniper Networks, Inc.

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, as per *Multiprotocol Interconnect over Frame Relay* (RFC 2427 [1490]).

ether-vpls-over-ppp—For E1, T1, E3, T3 and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over PPP encapsulation to support Bridged Ethernet over PPP encapsulated TDM interfaces for VPLS applications.

ethernet—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

extended-vlan-vpls—Use extended virtual LAN (VLAN) VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901.

NOTE: **The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.**

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only. J Series routers do not support **frame-relay-ppp** encapsulation.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting unlike media. When you use this encapsulation type, you can configure the **tcc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with **flexible-frame-relay** encapsulation.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect unlike media. The physical interface must be configured with **flexible-frame-relay** encapsulation.

Copyright © 2013, Juniper Networks, Inc.

multilink-frame-relay-end-to-end—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services,and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—For underlying Ethernet interfaces on J Series Services routers, use PPP over Ethernet encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface. You also use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 Series routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

ppp-over-ether-over-atm-llc—For underlying ATM interfaces on J Series Services routers only, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

vlan-bridge—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q tagging, flexible ethernet services, and bridging enabled, and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

vlan-ccc—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-vpls—Use Ethernet VLAN encapsulation on VPLS circuits.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution*
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

encapsulation (Logical Interface)

Syntax

encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-tcc-vc-mux | atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | ethernet-ccc | ethernet-vpls | ethernet-vpls-fr | frame-relay-ccc | frame-relay-ether-type | frame-relay-ether-type-tcc | frame-relay-ppp | frame-relay-tcc | gre-fragmentation | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls);

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces rlsq *number* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (**vlan-ccc** and **vlan-tcc** options only).
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access Routers. Only the **atm-ccc-cell-relay** and **atm-ccc-vc-mux** options are supported on ACX Series routers.

Description

Configure a logical link-layer encapsulation type.

Options

atm-ccc-cell-relay—Use ATM cell-relay encapsulation.

atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

atm-cisco-nlpid—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink Point-to-Point (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-ppp-llc—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over AAL5 LLC encapsulation.

atm-ppp-vc-mux—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over ATM AAL5 multiplex encapsulation.

atm-snap—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.

atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

Copyright © 2013, Juniper Networks, Inc.

atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

atm-vc-mux—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

ether-over-atm-llc—(All IP interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) For interfaces that carry IP traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, per RFC 2427, *Multiprotocol Interconnect over Frame Relay*.

NOTE: **The SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and the DS3/E3 MIC do not support Ethernet over Frame Relay encapsulation.**

ether-vpls-over-ppp—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Point-to-Point Protocol (PPP) encapsulation to support Bridged Ethernet over PPP-encapsulated TDM interfaces for VPLS applications.

ethernet—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

NOTE: **The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.**

Copyright © 2013, Juniper Networks, Inc.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE device over a time-division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only. J Series routers do not support frame-relay-ppp encapsulation.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

gre-fragmentation—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the do not fragment (DF) bit in the packet header. If the packet's size exceeds the tunnel's maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

multilink-frame-relay-end-to-end—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—For underlying Ethernet interfaces on J Series routers, use PPP over Ethernet encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface. You also use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

ppp-over-ether-over-atm-llc—(J Series routers and MX Series routers with MPCs using the ATM MIC with SFP only) For underlying ATM interfaces, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

vlan-bridge—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q tagging, flexible-ethernet-services, and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

vlan-ccc—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-vpls—Use Ethernet VLAN encapsulation on VPLS circuits.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring the Encapsulation for Layer 2 Switching TCCs*
- *Configuring Interface Encapsulation on Logical Interfaces*
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Circuit and Translational Cross-Connects Overview*
- *Identifying the Access Concentrator*
- *Configuring ATM Interface Encapsulation*
- *Configuring VLAN Encapsulation*
- *Configuring Extended VLAN Encapsulation*
- *Configuring ISDN Logical Interface Properties*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- Configuring ATM for Subscriber Access on page 587
- *Junos OS Services Interfaces Library for Routing Devices*
- *CoS on ATM IMA Pseudowire Interfaces Overview*
- *Configuring Policing on an ATM IMA Pseudowire*

Copyright © 2013, Juniper Networks, Inc.

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

Syntax enforce-strict-scale-limit-license;

Hierarchy Level [edit system services subscriber-management]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure the router to strictly enforce the subscriber scaling license, and to not allow the normal grace period. No additional subscribers are allowed to log in after the number of subscribers reaches the maximum allowed for the license.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Router to Strictly Enforce the Subscriber Scaling License on page 235

enhanced-policer

Syntax enhanced-policer

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 12.3.

Description Display the collection of policer detailed statistics. An FPC restart is required after changing this configuration.

When you commit a configuration that contains the **enhanced-policer** statement at the **[edit chassis]** hierarchy level, a warning message is displayed stating that all the FPCs in the router need to be rebooted for the configuration changes to become effective. At this point, you must confirm that you want to proceed with the reboot of the FPCs. If you do not reboot the FPCs, the FPCs return all 0s (zeros) when you perform a query for the retrieval of detailed statistics—for example, when you issue the **show firewall detail** command.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Router Chassis Configuration Statements*

equals (DHCP Relay Agent Option)

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.

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 323

Copyright © 2013, Juniper Networks, Inc.

entity-type

Syntax entity-type (host | mobility-agent);

Hierarchy Level [edit logical-systems *logical-system-name* services mobile-ip peer spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer spi *hexadecimal-value*],
[edit routing-instances *routing-instances-name* services mobile-ip peer spi *hexadecimal-value*],
[edit services mobile-ip peer spi *hexadecimal-value*]

Release Information Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Configure the security parameter for the peer entity—, either a mobile node, home agent, or foreign agent.

Options **host**—Use the mobile node in home agent

mobility-agent—Use the home agent or foreign agent

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

equals (Dynamic Profile)

Syntax equals *expression*;

Hierarchy Level [edit dynamic-profiles *profile-name* variables *variable-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure an expression for a user-defined variable that is evaluated at run time and returned as the variable value.

Options ***expression***—Expression evaluated to return a value for the user-defined variable.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Variable Expressions Overview on page 667
- Configuring User-Defined CoS Variables in a Dynamic Service Profile on page 1000

Copyright © 2013, 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 47
- Configuring RADIUS Server Parameters for Subscriber Access on page 36

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 © 2013, Juniper Networks, Inc.

excess-burst-size

Syntax excess-burst-size *bytes*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall three-color-policer *name* single-rate],
[edit firewall three-color-policer *policer-name* single-rate]

Release Information Statement introduced in Junos OS Release 7.4.
Support at the **[edit dynamic-profiles ... single-rate]** hierarchy level introduced in Junos Release OS 11.4.

Description For a single-rate three-color policer, configure the excess burst size (EBS) as a number of bytes. The EBS allows for moderate periods of bursting traffic that exceeds both the committed information rate (CIR) and the committed burst size (CBS).

NOTE: When you include the **excess-burst-size** statement in the configuration, you must also include the **committed-burst-size** and **committed-information-rate** statements at the same hierarchy level.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policing uses a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red based on the **excess-burst-size** statement included in the policer configuration.

During periods of traffic that conforms to the CIR, any unused portion of the guaranteed bandwidth capacity accumulates in the first token bucket, up to the maximum number of bytes defined by the CBS. If any accumulated bandwidth capacity overflows the first bucket, the excess accumulates in a second token bucket, up to the maximum number of bytes defined by the EBS.

A nonconforming traffic flow is categorized yellow if its size conforms to bandwidth capacity accumulated in the first token bucket. Packets in a yellow flow are marked with **medium-high** packet loss priority (PLP) and then passed through the interface.

A nonconforming traffic flow is categorized red if its size exceeds the bandwidth capacity accumulated in the second token bucket. Packets in a red traffic flow are marked with **high** PLP and then either passed through the interface or optionally discarded.

Options *bytes*—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 100,000,000,000 bytes

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Three-Color Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Dual Token Bucket Algorithms*
- *Determining Proper Burst Size for Traffic Policers*
- committed-burst-size on page 1516
- committed-information-rate on page 1518

excess-priority (Dynamic Schedulers)

Syntax

excess-priority (low | high | $junos-cos-scheduler-excess-priority | none);

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]

Release Information

Statement introduced in Junos OS Release 10.2.
Option **none** introduced in Junos OS Release 11.4.

Description

Determine the priority of excess bandwidth traffic on a scheduler in a dynamic profile.

Options

low—Excess traffic for this scheduler has low priority.

high—Excess traffic for this scheduler has high priority.

$junos-cos-scheduler-excess-priority—Variable for the excess-priority that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

none—System does not demote the priority of guaranteed traffic when the bandwidth exceeds the shaping rate or the guaranteed rate.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 1090
- scheduler on page 1968

Copyright © 2013, Juniper Networks, Inc.

excess-rate (Dynamic Schedulers)

Syntax excess-rate percent (*percentage* | $junos-cos-scheduler-excess-rate);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Determine the percentage of excess bandwidth traffic to share.

Options ***percentage***—Percentage of the excess bandwidth to share.

Range: 0 through 100 percent

$junos-cos-scheduler-excess-rate—Variable for the excess rate that is specified for a scheduler. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 1090
- output-traffic-control-profile on page 1839

excess-rate (Dynamic Traffic Shaping)

Syntax excess-rate (percent *percentage* | $junos-cos-excess-rate) | proportion *value*);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description For an MPC interface, determine the percentage or proportion of excess bandwidth traffic to share for all priorities of traffic.

Options ***percentage***—Percentage of the excess bandwidth to share.

Range: 0 through 100 percent

value—Proportion of the excess bandwidth to share.

Range: 0 through 1000

$junos-cos-excess-rate—Variable for the excess rate that is specified for the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 1090
- **output-traffic-control-profile on page 1839**

 Copyright © 2013, Juniper Networks, Inc.

excess-rate-high (Dynamic Traffic Shaping)

Syntax excess-rate-high ((percent *percentage* | $junos-cos-excess-rate-high) | proportion *value*);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description For an MPC/MIC interface, determine the percentage of excess bandwidth for high-priority traffic to share.

Options ***percentage***—Percentage of the excess bandwidth to share.

Range: 0 through 100 percent

value—Proportion of the excess bandwidth to share.

Range: 0 through 1000

$junos-cos-excess-rate-high—Variable for the excess rate that is specified for high-priority traffic on the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 1090
- **output-traffic-control-profile on page 1839**

excess-rate-low (Dynamic Traffic Shaping)

Syntax excess-rate-low ((percent *percentage* | $junos-cos-excess-rate-low) | proportion *value*);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description For an MPC/MIC interface, determine the percentage of excess bandwidth for low-priority traffic to share.

Options ***percentage***—Percentage of the excess bandwidth to share.

Range: 0 through 100 percent

value—Proportion of the excess bandwidth to share.

Range: 0 through 1000

$junos-cos-excess-rate-low—Variable for the excess rate that is specified for low-priority traffic on the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 1090
- **output-traffic-control-profile on page 1839**

 Copyright © 2013, 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 ];
    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 ];
    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 ];
    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 ];
    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 ];
    max-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    max-data-rate-up [ access-request | accounting-start | accounting-stop ];
    max-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
    max-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
    min-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    min-data-rate-up [ access-request | accounting-start | accounting-stop ];
    min-lp-data-rate-dn [ access-request | accounting-start | accounting-stop ];
    min-lp-data-rate-up [ access-request | accounting-start | accounting-stop ];
    nas-identifier [ access-request | accounting-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 ];
```

```
    upstream-calculated-qos-rate [ 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** introduced in Junos OS Release 11.4.
Options **cos-shaping-rate** and **filter-id** introduced in Junos OS Release 13.2.

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.

 Copyright © 2013, 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.
- **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.
- **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.
- **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.
- **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.

Copyright © 2013, Juniper Networks, Inc.

- **input-gigawords**—RADIUS attribute 52, Acct-Input-Gigawords.
- **interface-description**—Juniper Networks VSA 26-53, Interface-Desc.
- **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 25-43, Acct-Output-Gigapackets.
- **output-gigawords**—RADIUS attribute 53, Acct-Output-Gigawords.
- **upstream-calculated-qos-rate**—Juniper Networks VSA 26-142

RADIUS message type

- **access-request**—RADIUS Access-Accept 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 RADIUS Server Parameters for Subscriber Access on page 36

Copyright © 2013, Juniper Networks, Inc.

exclude (Dynamic MLD Interface)

Syntax exclude;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the group to operate in exclude mode on the dynamic interface. In exclude mode all sources except the address configured are accepted for the group. By default, the group operates in include mode.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation
- *Enabling MLD Static Group Membership*

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 215
- Extended DHCP Local Server Overview on page 200
- Address-Assignment Pools Overview on page 167

fail-over-within-preference (L2TP LAC)

Syntax fail-over-within-preference;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 10.4.

Description Enable L2TP LAC tunnel selection within a preference level. When the router is unable to connect to a destination at a given preference level, it attempts to connect to another destination at the same level. By default, when a connection attempt fails at one preference level, the next attempt is made at the next lower level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring LAC Tunnel Selection Failover Within a Preference Level on page 406
- Configuring the L2TP LAC Tunnel Selection Parameters on page 406

 Copyright © 2013, Juniper Networks, Inc.

failure-action

Syntax failure-action (clear-binding | clear-binding-if-interface-up | log-only);

Hierarchy Level [edit system services dhcp-local-server liveness-detection],
[edit system services dhcp-local-server dhcpv6 liveness-detection],
[edit forwarding-options dhcp-relay liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection],
[edit system services dhcp-local-server group *group-name* liveness-detection],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the action the router (or switch) takes when a liveness detection failure occurs.

Options

clear-binding—The client session is cleared when a liveness detection failure occurs.

clear-binding-if-interface-up—The client session is cleared only when a liveness detection failure occurs and the local interface is detected as being up.

log-only—A message is logged to indicate the event; no action is taken and DHCP is left to manage the failure.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- DHCP Liveness Detection Overview on page 250
- Configuring Detection of DHCP Local Server Client Connectivity on page 251
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 318
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

Copyright © 2013, 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.

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 167
- Configuring Address-Assignment Pools on page 168

Copyright © 2013, Juniper Networks, Inc.

family (Dynamic Firewalls)

Syntax

```
family family {
  fast-update-filter filter-name {
    interface-specific;
    match-order [match-order];
    term term-name {
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
      only-at-create;
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Configure protocol family information for firewall filters in a dynamic profile.

Options

family—Protocol family:

- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Fast Update Filters on page 1183

family (Dynamic Demux Interface)

Syntax

```
family family {
  access-concentrator name;
  address address;
  demux-source {
    source-address;
  }
  duplicate-protection;
  dynamic-profile profile-name;
  filter {
    input filter-name;
    output filter-name;
  }
  mac-validate (loose | strict);
  max-sessions number;
  max-sessions-vsa-ignore;
  service-name-table table-name;
  short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
  unnumbered-address interface-name <preferred-source-address address>;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 9.3.
Option **pppoe** introduced in Junos OS Release 11.2.

Description

Configure protocol family information for the logical interface.

NOTE: Not all subordinate stanzas are available to every protocol family.

Options

family—Protocol family:

- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite
- **pppoe**—(MX Series routers with MPCs only) Point-to-Point Protocol over Ethernet

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- For information about static IP demux interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

family (Dynamic PPPoE)

Syntax

```
family family {
  unnumbered-address interface-name;
  address address;
  service {
    input {
      service-set service-set-name {
        service-filter filter-name;
      }
      post-service-filter filter-name;
    }
    output {
      service-set service-set-name {
        service-filter filter-name;
      }
    }
  }
  filter {
    input filter-name {
      precedence precedence;
    }
    output filter-name {
      precedence precedence;
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit"]

Release Information

Statement introduced in Junos OS Release 10.1.

Description

Configure protocol family information for the logical interface.

Options

family—Protocol family:

- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Basic PPPoE Dynamic Profile on page 900
- Configuring a PPPoE Dynamic Profile with Additional Options on page 903
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

family (Dynamic Standard Interface)

Syntax

```
family family {
  access-concentrator name;
  address address;
  duplicate-protection;
  dynamic-profile profile-name;
  filter {
    adf {
      counter;
      input-precedence precedence;
      not-mandatory;
      output-precedence precedence;
      rule rule-value;
    }
    input filter-name {
      precedence precedence;
    }
    output filter-name {
      precedence precedence;
    }
  }
  mac-validate (loose | strict);
  max-sessions number;
  max-sessions-vsa-ignore;
  rpf-check {
    fail-filter filter-name;
    mode loose;
  }
  service {
    input {
      service-set service-set-name {
        service-filter filter-name;
      }
      post-service-filter filter-name;
    }
    output {
      service-set service-set-name {
        service-filter filter-name;
      }
    }
  }
  service-name-table table-name
  short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
  unnumbered-address interface-name <preferred-source-address address>;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 9.2.
Option **pppoe** introduced in Junos OS Release 11.2.

Description

Configure protocol family information for the logical interface.

Copyright © 2013, Juniper Networks, Inc.

NOTE: **Not all subordinate stanzas are available to every protocol family.**

Options

family—Protocol family:

- **inet**—IP version 4 suite
- **inet6**—IP version 6 suite
- **pppoe**—(MX Series routers with MPCs only) Point-to-Point Protocol over Ethernet
- **vpls**—Virtual private LAN service

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

- For general information about configuring static interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.
- "Configuring the Protocol Family," in the *Junos OS Network Interfaces Library for Routing Devices*.

fast-update-filter (Dynamic Firewalls)

Syntax

```
fast-update-filter filter-name {
  interface-specific;
  match-order [match-order];
  term term-name {
    from {
      match-conditions;
    }
    then {
      action;
      action-modifiers;
    }
    only-at-create;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall family *family*]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Configure fast update firewall filters in a dynamic profile.

Options

filter-name—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

The statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Fast Update Filters on page 1183

Copyright © 2013, Juniper Networks, Inc.

filter (Configuring)

Syntax

```
filter filter-name {
   accounting-profile name;
   enhanced-mode;
   interface-shared;
   interface-specific;
   physical-interface-filter;
   term term-name {
      ... term configuration ...
   }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall family *family-name*],
[edit firewall family *family-name*],
[edit logical-systems *logical-system-name* firewall family *family-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Logical systems support introduced in Junos OS Release 9.3.
physical-interface-filter statement introduced in Junos OS Release 9.6.
Support at the **[edit dynamic-profiles ... family family-name]** hierarchy level introduced in Junos OS Release 11.4.
Support for the **interface-shared**> statement introduced in Junos OS Release 12.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure firewall filters.

Options

filter-name—Name that identifies the filter. This must be a non-reserved string of not more than 64 characters. To include spaces in the name, enclose it in quotation marks (“ ”). In Junos OS Release 9.0 and later, you can no longer use special characters within the name of a firewall filter. Firewall filter names are restricted from having the form __.*__ (beginning and ending with underscores) or __.* (beginning with an underscore.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Guidelines for Configuring Firewall Filters*
- *Guidelines for Applying Firewall Filters*
- *Configuring Multifield Classifiers*
- *Using Multifield Classifiers to Set PLP*
- *simple-filter (Configuring)*

Copyright © 2013, Juniper Networks, Inc.

filter (Dynamic Firewalls)

Syntax

```
filter {
  adf {
    counter;
    input-precedence precedence;
    not-mandatory;
    output-precedence precedence;
    rule rule-value;
  }
  input filter-name {
    precedence precedence;
    shared-name filter-shared-name;
  }
  output filter-name {
    precedence precedence;
    shared-name filter-shared-name;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos–interface–unit" family *family*]

Release Information

Statement introduced in Junos OS Release 9.2.
Support at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family*]** hierarchy level introduced in Junos OS Release 10.1.
shared-name statement added in Junos OS Release 12.2.

Description

Apply a dynamic filter to an interface. You can configure filters for either **family inet** or **family inet6**, and the filters can be classic filters, fast update filters, or (for the **adf** statement) Ascend-Data-Filters. Only the Internet Protocol version 4 (IPv4) protocol family is currently supported for dynamic PPPoE logical interfaces.

Options

input *filter-name*—Name of one filter to evaluate when packets are received on the interface.

output *filter-name*—Name of one filter to evaluate when packets are transmitted on the interface.

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

- For general information about configuring firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*
- Dynamic Firewall Filters Overview on page 1132

Copyright © 2013, Juniper Networks, Inc.

- Classic Filters Overview on page 1133
- Basic Classic Filter Syntax on page 1135

filter (Dynamic Interface Unit)

Syntax

```
filter {
   input filter-name;
   output filter-name;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number*],

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Apply a dynamic filter to an interface, regardless of its family type.

Options

input *filter-name*—Name of one filter to evaluate when packets are received on the interface.

output *filter-name*—Name of one filter to evaluate when packets are transmitted on the interface.

The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- For general information about configuring firewall filters, see the *Routing Policy Feature Guide for Routing Devices*
- Dynamic Firewall Filters Overview on page 1132
- Classic Filters Overview on page 1133
- Basic Classic Filter Syntax on page 1135
- Dynamically Attaching Statically Created Filters for Any Interface Type on page 1170

Copyright © 2013, Juniper Networks, Inc.

filter-specific

Syntax filter-specific;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall policer *policer-name*],
[edit firewall family inet prefix-action *name*],
[edit firewall policer *policer-name*],
[edit logical-systems *logical-system-name* firewall policer *policer-name*],
[edit logical-systems *logical-system-name* firewall family inet prefix-action *name*]

Release Information Statement introduced before Junos OS Release 7.4.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... policer *policer-name*]** hierarchy level introduced in Junos OS Release 11.4.

Description Set the prefix-specific action or policer to operate in *filter-specific* mode, meaning that a single policer and counter are shared by all filter terms that reference the prefix-specific action or policer. By default, the prefix-specific action or policer operates in *term-specific* mode, meaning that a separate policer and counter are used for each filter term that references the prefix-specific action or policer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Filter-Specific Policer Overview*
- *Prefix-Specific Counting and Policing Overview*
- *Filter-Specific Counter and Policer Set Overview*

 Copyright © 2013, Juniper Networks, Inc.

firewall (Dynamic Firewalls)

Syntax

```
firewall {
  family family {
    fast-update-filter filter-name {
      interface-specific;
      match-order [match-order];
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
        only-at-create;
      }
    }
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure firewall filters in a dynamic profile.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Fast Update Filters on page 1183

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 323

Copyright © 2013, Juniper Networks, Inc.

forward-snooped-clients (DHCP Local Server)

Syntax

forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server],
[edit logical-systems *logical-system-name* system services dhcp-local-server],
[edit routing-instances *routing-instance-name* system services dhcp-local-server],
[edit system services dhcp-local-server]

Release Information

Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure how the DHCP local server handles DHCP snooped packets on specific interfaces.

Options

all-interfaces—Perform the action on all interfaces.

configured-interfaces—Perform the action only on configured interfaces.

non-configured-interfaces—Perform the action only on nonconfigured interfaces.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- DHCP Snooping Support on page 299
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 223

forward-snooped-clients (DHCP Relay Agent)

Syntax

forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);

Hierarchy Level

[edit forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay]

Release Information

Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure how DHCP relay agent handles DHCP snooped packets on specific interfaces. The router or switch determines the DHCP snooping action to perform based on a combination of the **forward-snooped-clients** configuration and the configuration of either the **allow-snooped-clients** statement or the **no-allow-snooped-clients** statement.

The router (or switch) also uses this statement to determine how to handle snooped BOOTREPLY packets received on nonconfigured interfaces.

Options

all-interfaces—Perform the action on all interfaces.

configured-interfaces—Perform the action only on configured interfaces.

non-configured-interfaces—Perform the action only on nonconfigured interfaces.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- DHCP Snooping Support on page 299
- Configuring DHCP Snooping for DHCP Relay Agent on page 301

Copyright © 2013, 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.

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 471
- Configuring Diameter Network Elements on page 473

forwarding-class (Dynamic Scheduler Maps)

Syntax forwarding-class *class-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service scheduler-maps *map-name*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Associate a scheduler with a scheduler map.

Options *class-name*—Name of the forwarding class.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973

forwarding-class (Subscriber Secure Policy)

Syntax forwarding-class *class-name*;

Hierarchy Level [edit services radius-flow-tap]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify forwarding class that is applied to mirrored packets sent to a mediation device.

Options *class-name*—Name of the forwarding class.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

 Copyright © 2013, Juniper Networks, Inc.

fpc (MX Series 3D Universal Edge Routers)

Syntax

```
fpc slot-number {
  inline-services {
    flow-table-size {
      ipv4-flow-table-size units;
      ipv4-flow-table-size  units;
    }
  }
  pic number {
    inline-services {
      bandwidth (1g | 10g);
    }
    port-mirror-instance port-mirroring-instance-name-pic-level;
    tunnel-services {
      bandwidth (1g | 10g)
    }
  }
  port-mirror-instance port-mirroring-instance-name-fpc-level;
}
```

Hierarchy Level

[edit chassis]

Release Information

Statement introduced in Junos OS Release 8.2.
port-mirror-instance option introduced in Junos OS Release 9.3.

Description

Configure properties for the DPC or MPC and corresponding Packet Forwarding Engines to create tunnel interfaces.

(MX Series Virtual Chassis only) To configure properties for MPCs in a member router in an MX Series Virtual Chassis configuration, you must specify the router's Virtual Chassis member number *before* the **fpc** statement. Specify the member number in the form **member** ***member-id***, where ***member-id*** is 0 or 1. If you do not specify the member number before the **fpc** statement, the commit operation fails and the software displays an error message indicating that the **fpc** statement must include the member number for routers in Virtual Chassis mode.

Options

fpc ***slot-number***—Specify the slot number of the DPC.

Range: 0 through 11

pic ***number***—Specify the number of the Packet Forwarding Engine. Each DPC includes four Packet Forwarding Engines.

Range: 0 through 4

port-mirror-instance ***port-mirroring-instance-name-fpc-level***—Associate a port-mirroring instance with the DPC and its corresponding PICs. The port-mirroring instance is configured under the **[edit forwarding-options port-mirroring]** hierarchy level.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Port-Mirroring Instances on MX Series 3D Universal Edge Routers*
- Enabling Inline Service Interfaces on page 427

Copyright © 2013, Juniper Networks, Inc.

frame-mode (Dynamic Traffic Shaping)

Syntax

frame-mode (bytes | $junos-cos-byte-adjust | frame-mode-bytes *frame-mode-bytes* |$junos-cos-byte-adjust-frame);

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*],
[edit class-of-service traffic-control-profiles *profile-name*]

Release Information

Statement introduced in Junos OS Release 10.2.
Variable ***$junos-cos-byte-adjust-frame*** introduced in Junos OS Release 13.1.

Description

Configure the mode to shape downstream ATM traffic based as frames.

Default

The default is **frame-mode**.

Options

bytes—Byte adjustment value for the **cell-mode** or **frame-mode** shaping options.

$junos-cos-byte-adjust—Predefined variable for byte adjustment that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

frame-mode-bytes *frame-mode-bytes*—Overhead bytes when in frame-mode. Traffic shaping is based on the number of bytes in the frame, without regard to cell encapsulation or padding overhead.

$junos-cos-byte-adjust-frame—Predefined variable for frame mode shaping. This variable can not be used when the **overhead-accounting bytes *bytes*** option is configured.

BEST PRACTICE: We recommend using the **frame-mode-bytes *frame-mode-bytes*** option rather than the **bytes** option.

Range: –120 through 124 bytes

NOTE: If you specify a value for the **bytes *bytes*** option, you cannot specify a value for either the **frame-mode-bytes** option.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- CoS Adjustment Control Profiles Overview on page 1084
- Configuring CoS Adjustment Control Profiles on page 1106
- **adjustment-control-profiles on page 1442**
- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 1092

- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072
- *egress-shaping-overhead*
- bytes on page 1488
- cell-mode on page 1494

from

Syntax

```
from {
  application [junos-http, junos-https, junos-httpproxy];
  destination-address (CoS) address <except>;
  destination-prefix-list list-name <except>;
}
```

Hierarchy Level

[edit services (captive-portal-content-delivery) captive-portal-content-delivery rule *rule-name* term *term-name*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Specify input conditions for a captive portal term.

Options

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

- Redirecting HTTP Requests Overview on page 1225
- For information on match conditions, see the description of firewall filter match conditions in the *Routing Policy Feature Guide for Routing Devices*.

Copyright © 2013, Juniper Networks, Inc.

from (Subscriber Secure Policy)

Syntax

```
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;
}
```

Hierarchy Level [edit services radius-flow-tappolicy *policy-name* inet| inet6]

Release Information Statement introduced in Junos OS Release 12.3.

Description Define the match criteria for the drop-policy rule.

The remaining statements are explained separately.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

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.

Description Specify the application (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. Gx-Plus, JSRC, and packet-triggered subscribers are the applications currently supported.

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 471
- Configuring Diameter Network Elements on page 473

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.

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, 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 471
- Configuring Diameter Network Elements on page 473

gateway-name (Tunnel Profile Remote Gateway)

Syntax gateway-name *server-name*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id* remote-gateway]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the hostname expected by the remote gateway—the LNS—from the source gateway—the LAC—when you set up a tunnel.

Options ***server-name***—Name of the LNS.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 403

 Copyright © 2013, Juniper Networks, Inc.

gateway-name (Tunnel Profile Source Gateway)

Syntax gateway-name *client-name*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id* source-gateway]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the hostname provided by the source gateway—the LAC—to the remote gateway—the LNS—when you set up a tunnel.

Options ***client-name***—Name of the LAC.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

generic

Syntax generic;

Hierarchy Level [edit logical-systems *logical-system-name* services mobile-ip access-type],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* services mobile-ip access-type],
[edit routing-instances *routing-instance-name* services mobile-ip access-type],
[edit services mobile-ip access-type]

Release Information Statement introduced in Junos OS Release 9.5.

Description Disable WiMAX features for Mobile IP home agent, preventing interoperation in a WiMAX environment.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Access Type for Mobile IP on page 573

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.

Description Configure global attributes for the Gx-Plus application.

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

- Configuring Gx-Plus on page 549

grace-period

Syntax grace-period *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes]

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 168

 Copyright © 2013, Juniper Networks, Inc.

gres-route-flush-delay (Subscriber Management)

Syntax gres-route-flush-delay;

Hierarchy Level [edit system services subscriber-management]

Release Information Statement introduced in Junos OS Release 11.2.

Description For a subscriber network configured with either nonstop active routing (NSR) or graceful restart, configure the router to wait 180 seconds (3 minutes) before removing (flushing) static or dynamic access routes and access-internal routes from the forwarding table after a graceful Routing Engine switchover (GRES) has taken place.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 693
- Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 688

group (DHCP Local Server)

Syntax

```
group group-name {
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      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 {
    exclude;
    overrides {
      client-discover-match <option60-and-option82>;
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
      rapid-commit;
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        detection-time {
          threshold milliseconds;
        }
        session-mode(automatic | multihop | singlehop);
```


Copyright © 2013, Juniper Networks, Inc.

```
            holddown-interval milliseconds;
          }
        }
      }
      overrides {
        client-discover-match <option60-and-option82>;
        delegated-pool;
        interface-client-limit number;
        process-inform {
          pool pool-name;
        }
        rapid-commit;
      }
      reconfigure {
        attempts attempt-count;
        clear-on-abort;
        strict;
        timeout timeout-value;
        token token-value;
        trigger {
          radius-disconnect;
        }
      }
      route-suppression;
      service-profile dynamic-profile-name;
    }
```

Hierarchy Level

[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Configure a group of interfaces that have a common configuration, such as authentication parameters. A group must contain at least one interface.

Options

group-name—Name of the group.

The remaining statements are explained separately.

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 200
- Grouping Interfaces with Common DHCP Configurations on page 217
- Using External AAA Authentication Services with DHCP on page 214
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

Copyright © 2013, Juniper Networks, Inc.

group (DHCP Relay Agent)

Syntax

```
group group-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;
      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 {
    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 {
      ...
    }
    service-profile dynamic-profile-name;
    trace;
```

```
        upto upto-interface-name;
    }
    overrides {
        allow-snooped-clients;
        always-write-giaddr;
        always-write-option-82;
        client-discover-match <option60-and-option82>;
        disable-relay;
        interface-client-limit number;
        layer2-unicast-replies;
        no-allow-snooped-clients;
        no-bind-on-request;
        proxy-mode;
        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-option {
        option-number option-number;
        default-action {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
        equals (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
        starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
    }
    relay-option-82 {
        circuit-id {
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82;
        }
    }
    route-suppression;
    service-profile dynamic-profile-name;
}
```


Copyright © 2013, Juniper Networks, Inc.

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.

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 1548
- Extended DHCP Relay Agent Overview on page 276
- Group-Specific DHCP Relay Options on page 292
- Grouping Interfaces with Common DHCP Configurations on page 217
- Using External AAA Authentication Services with DHCP on page 214
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

group (Dynamic IGMP Interface)

Syntax

For group configuration with a source, use the following syntax:

```
group ip-address {
  source ip-address;
}
```

For group configuration without a source, use the following syntax:

```
group group;
```

Hierarchy Level

[edit dynamic-profiles *profile-name* protocols igmp interface *interface-name* static],

Release Information

Statement introduced in Junos OS Release 9.2.

Description

When configuring with a source address, configure the IGMP multicast group address that receives data on an interface and a source address for certain packets. For configuration without a source address, configure only the IGMP multicast group address that receives data on an interface.

Options

ip-address—Group IP address.

group—Name of group.

NOTE: **You must specify a unique address for each group.**

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about configuring static group membership, see "Enabling IGMP Static Group Membership" in the *Multicast Protocols Feature Guide for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

group (Dynamic MLD Interface)

Syntax

```
group multicast-group-address {
  exclude;
  group-count number;
  group-increment increment;
  source ip-address {
    source-count number;
    source-increment increment;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* protocols mld interface *interface-name* static]

Release Information

Statement introduced in Junos OS Release 10.1.

Description

The MLD multicast group address and (optionally) the source address for the multicast group being dynamically configured on an interface.

Options

multicast-group-address—Address of the group.

NOTE: **You must specify a unique address for each group.**

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Enabling MLD Static Group Membership*

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.

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 500
- Creating a Static Subscriber Group on page 504

Copyright © 2013, Juniper Networks, Inc.

group-count (Dynamic MLD Interface)

Syntax group-count *number*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the number of static groups to be created over the dynamic interface.

Options *number*—Number of static groups.

Default: 1

Range: 1 through 512

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Enabling MLD Static Group Membership*

group-increment (Dynamic MLD Interface)

Syntax group-increment *increment*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the number of times the address should be incremented for each static group created on a dynamic interface. The increment is specified in a format similar to an IPv6 address.

Options *increment*—Number of times the address should be incremented.

Default: ::1

Range: ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Enabling MLD Static Group Membership*

group-limit (Dynamic IGMP Interface)

Syntax group-limit *limit*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*],

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on a dynamic logical interface. After this limit is reached, new reports will be ignored and all related flows are not flooded on the logical interface.

Default By default, there is no limit to the number of multicast groups that can join the interface.

Options ***limit***—group limit value for the interface.

Range: 1 through 32767

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about limiting the number of multicast group joins for an IGMP logical interface, see "*Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces*" in the *Multicast Protocols Feature Guide for Routing Devices*

 Copyright © 2013, Juniper Networks, Inc.

group-limit (Dynamic MLD Interface)

Syntax group-limit *limit*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure a limit for the number of multicast groups (or [S,G] channels in MLDv2) allowed on a dynamic logical interface. After this limit is reached, new reports will be ignored and all related flows are not flooded on the logical interface.

Default By default, there is no limit to the number of multicast groups that can join the interface.

Options ***limit***—group limit value for the interface.

Range: 1 through 32767

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about limiting the number of multicast group joins for an MLD logical interface, see "*Configuring the Number of MLD Multicast Group Joins on Logical Interfaces*" in the *Multicast Protocols Feature Guide for Routing Devices*

group-policy (Dynamic IGMP Interface)

Syntax group-policy *policy-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Compare the IGMPv2 or IGMPv3 group against the specified group policy, after receiving an IGMP report, and perform the action configured in that policy (for example, reject the report).

Options ***policy-name***—Name of the group policy.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about rejecting unwanted reports for an IGMP interface, see "Filtering Unwanted IGMP Reports at the IGMP Interface Level" in the *Multicast Protocols Feature Guide for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

group-policy (Dynamic MLD Interface)

Syntax group-policy *policy-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Compare the MLDv1 or MLDv2 group against the specified group policy, after receiving an MLD report, and perform the action configured in that policy (for example, reject the report).

Options ***policy-name***—Name of the group policy.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Filtering Unwanted MLD Reports at the MLD Interface Level*

 Copyright © 2013, Juniper Networks, Inc.

group-profile (Group Profile)

Syntax

```
group-profile profile-name {
  l2tp {
    interface-id interface-id;
    lcp-renegotiation;
    local-chap;
    maximum-sessions-per-tunnel number;
  }
  ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-pool pool-id;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    ppp-options {
      chap;
      pap;
    }
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the group profile.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options ***profile-name***—Name assigned to the group profile.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 418

guaranteed-rate (Dynamic Traffic Shaping)

Syntax

guaranteed-rate (*rate* | $junos-cos-guaranteed-rate) <burst-size [*bytes* | $junos-cos-guaranteed-rate-burst]>;

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.2.
The **$junos-cos-guaranteed-rate** variable introduced in Junos OS Release 9.4.
Option **burst-size** introduced in Junos OS Release 11.4.

Description

Configure a guaranteed minimum rate for a logical interface.

Default

If you do not include this statement and you do not include the **delay-buffer-rate** statement, the logical interface receives a minimal delay-buffer rate and minimal bandwidth equal to 2 MTU-sized packets.

Options

rate—Guaranteed rate in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1000 through 160,000,000,000 bps

$junos-cos-guaranteed-rate—Junos predefined variable that is replaced with the guaranteed rate obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

burst-size *bytes*—(Optional) Maximum burst size, in bytes.

Range: 0 through 1,000,000,000

$junos-cos-guaranteed-rate-burst—(Optional) Variable for the burst-size that is specified for the guaranteed rate. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profile]** hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Traffic Scheduling and Shaping for Subscriber Access on page 971
- **output-traffic-control-profile on page 1839**

Copyright © 2013, 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.

Description Configure the Gx-Plus application to interact with a PCRF to authorize and provision subscribers.

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

- Configuring Gx-Plus on page 549

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 168

hello-interval

Syntax hello-interval *seconds*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the keepalive timer for L2TP tunnels.

NOTE: **Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.**

Options ***seconds***—Interval, in seconds, after which the server sends a hello message if no messages are received. A value of **0** means that no hello messages are sent.

Default: 60 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Timers for L2TP Tunnels*
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429

 Copyright © 2013, Juniper Networks, Inc.

hierarchical-policer

Syntax

```
hierarchical-policer hierarchical-policer-name {
  aggregate {
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      discard;
    }
  }
  premium {
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      discard;
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall],
[edit firewall]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... firewall]** hierarchy level introduced in Junos OS Release 11.4.

Description

On M40e, M120, and M320 edge routers with Flexible PIC Concentrator (FPC) input as FFPC and FPC output as SFPC, and on MX Series, T320, T640, and T1600 edge routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify a hierarchical policer.

Options

hierarchical-policer-name—Name that identifies the policer. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Hierarchical Policer Configuration Overview*
- *Hierarchical Policers*
- aggregate (Hierarchical Policer) on page 1447
- bandwidth-limit (Hierarchical Policer) on page 1475
- burst-size-limit (Hierarchical Policer) on page 1484

Copyright © 2013, Juniper Networks, Inc.

- if-exceeding (Hierarchical Policer) on page 1692
- premium (Hierarchical Policer) on page 1894

Copyright © 2013, Juniper Networks, Inc.

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)

Syntax

```
hierarchical-scheduler {
   implicit-hierarchy;
   maximum–hierarchy–levels number;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.
implicit-hierarchy option introduced in Junos OS Release 13.1.

Description On MX Series routers with MPC/MICs, configure the parameters for hierarchical scheduling on the interface.

NOTE: To enable hierarchical scheduling on MX Series routers, configure the **hierarchical-scheduler** statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.

Options

implicit-hierarchy—Configure three-level hierarchical scheduling. When you include the **implicit-hierarchy** option, a hierarchical relationship is formed between the CoS scheduler nodes at level 1, level 2, and level 3. The **implicit-hierarchy** option is supported only on MPC/MIC subscriber interfaces and interface sets running over aggregated Ethernet on MX Series routers.

maximum-hierarchy-levels ***number***—Configure two-level hierarchical scheduling. Specify the maximum number of hierarchical scheduling levels allowed for node scaling. The only supported value is 2. The **maximum-hierarchy-levels** option is supported on MPC/MIC or EQ DPC subscriber interfaces and interface sets running over aggregated Ethernet on MX Series routers.

- If you include the **maximum-hierarchy-levels** option, interface sets are allowed only at level 3; they are not allowed at level 2. In this case, if you configure a level 2 interface set, you generate Packet Forwarding Engine errors.
- If you do not include the **maximum-hierarchy-levels** option, interface sets can be at either level 2 or level 3, depending on whether the member logical interfaces within the interface set have a traffic control profile. If any member logical interface has a traffic control profile, then the interface set is a level 2 CoS scheduler node. If no member logical interface has a traffic control profile, the interface set is at level 3.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers on page 956
- Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links on page 1023
- *Junos OS Class of Service Library for Routing Devices*
- *Configuring Hierarchical Schedulers for CoS*
- Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview on page 925

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 167
- Configuring Address-Assignment Pool Usage Threshold Traps on page 172

 Copyright © 2013, 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 167
- Configuring Address-Assignment Pool Usage Threshold Traps on page 172

holddown-interval

Syntax holddown-interval *milliseconds*;

Hierarchy Level [edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the time (in milliseconds) for which Bidirectional Forwarding Detection (BFD) holds a session up notification.

Options *milliseconds*—Interval specifying how long a BFD session must remain up before a state change notification is sent.

Range: 0 through 255,000

Default: 0

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

 Copyright © 2013, Juniper Networks, Inc.

home-agent (Mobile IP Dynamic Assignment)

Syntax

```
home-agent {
  nai (name@domain | @domain) {
    home-agent ip-address;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip dynamic-home-assignment],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip dynamic-home-assignment],
[edit routing-instances *routing-instances-name* services mobile-ip dynamic-home-assignment],
[edit services mobile-ip dynamic-home-assignment]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* services mobile-ip** dynamic-home-assignment**]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip** dynamic-home-assignment**]**, and **[edit routing-instances *routing-instances-name* services mobile-ip** dynamic-home-assignment**]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Configure the IP address to which registration requests are sent as part of the home agent's dynamic assignment rule.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring Dynamic Home Assignment for the Mobile Node on page 572

home-agent (Mobile IP Network Address Identifier)

Syntax

home-agent *ip-address*;

Hierarchy Level

[edit services mobile-ip dynamic-home-assignment home-agent nai *name*@*domain*],
[edit logical-systems *logical-system-name* services mobile-ip dynamic-home-assignment home-agent nai *name*@*domain*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* services mobile-ip dynamic-home-assignment home-agent nai *name*@*domain*],
[edit routing-instances *routing-instance-name* services mobile-ip dynamic-home-assignment home-agent nai *name*@*domain*],
[edit services mobile-ip dynamic-home-assignment home-agent nai @*domain*],
[edit logical-systems *logical-system-name* services mobile-ip dynamic-home-assignment home-agent nai @*domain*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* services mobile-ip dynamic-home-assignment home-agent nai @*domain]*,
[edit routing-instances *routing-instance-name* services mobile-ip dynamic-home-assignment home-agent nai @*domain*]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure the IP address to which registration requests are sent as part of the home agent's dynamic assignment rule.

Options

ip-address—IP address of the home agent

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring Dynamic Home Assignment for the Mobile Node on page 572

Copyright © 2013, Juniper Networks, Inc.

home-agent (Mobile IP Networks)

Syntax

```
home-agent {
  enable-service interface-name;
  virtual-network {
    home-agent-address ip-address {
      registration-lifetime seconds;
      revocation-required;
      timestamp-tolerance seconds;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip],
[edit routing-instances *routing-instances-name* services mobile-ip],
[edit services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* services mobile-ip]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip]**, and **[edit routing-instances *routing-instances-name* services mobile-ip]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the virtual networks and non-virtual networks for the Mobile IP home agent.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569

home-agent-address

Syntax

```
home-agent-address ip-address {
   registration-lifetime seconds;
   revocation-required;
   timestamp-tolerance seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip home-agent virtual-network],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network],
[edit services mobile-ip home-agent virtual-network]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **edit logical-systems *logical-system-name* services mobile-ip home-agent** virtual-network], **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent** virtual-network], and **[edit routing-instances *routing-instances-name* services mobile-ip home-agent** virtual-network] hierarchy levels introduced in Junos OS Release 9.5.

Description

Defines addressing for the virtual network of the Mobile IP home agent.

Options

ip-address—For virtual networks, the loopback IP address for the virtual network. For non-virtual networks, a public address.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

Copyright © 2013, 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.

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 167
- Configuring Address-Assignment Pools on page 168

host (Diameter Origin)

Syntax host *hostname*;

Hierarchy Level [edit diameter origin]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the name of the host that originates the Diameter message.

Options ***hostname***—Name of the message origin host. Supplied as the value of Origin-Host AVP for all messages sent by 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 471
- Configuring the Origin Attributes of the Diameter Instance on page 472

identification (Tunnel Profile)

Syntax identification *name*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the assignment ID of an L2TP tunnel. L2TP sessions with the same tunnel assignment identification and destination are grouped into the same tunnel.

Options ***name***—Tunnel assignment ID; string of up to 32 alphanumeric characters.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

 Copyright © 2013, Juniper Networks, Inc.

idle-timeout (Access)

Syntax idle-timeout *seconds*;

Hierarchy Level [edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the idle timeout for a user. The router might consider a PPP session to be idle because of the following reasons:

- There is no ingress traffic on the PPP session.
- There is no egress traffic.
- There is neither ingress or egress traffic on the PPP session.
- There is no ingress or egress PPP control traffic. This is applicable only if keepalives are enabled.

Options ***seconds***—Number of seconds a user can remain idle before the session is terminated.

Range: 0 through 4,294,967,295 seconds

Default: 0

NOTE: The **[edit access]** hierarchy is not available on QFabric systems.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring PPP Properties for a Client-Specific Profile*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 418

idle-timeout (L2TP)

Syntax idle-timeout *seconds*;

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify how long a tunnel is active after its last session is terminated. The timer starts when the session is terminated and the tunnel is disconnected when the timer expires.

BEST PRACTICE: **Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing no services l2tp tunnel idle-timeout.**

Options *seconds*—Length of the idle timeout. A value of **0** creates a persistent tunnel; that is, the tunnel remains active indefinitely until the remote peer disconnects it or you issue the **clear services l2tp tunnel** command.

Range: 0 through 86,400

Default: 60

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting the L2TP Tunnel Idle Timeout on page 412
- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

 Copyright © 2013, Juniper Networks, Inc.

ieee-802.1 (Dynamic Classifiers)

Syntax ieee-802.1 (*classifier-name* | default) vlan-tag (inner | outer);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers]

Release Information Statement introduced in Junos OS Release 10.1.

Description Apply an IEEE-802.1 classifier to a subscriber interface in a dynamic profile.

Options *classifier-name*—Name of a **classifier** mapping configured at the **[edit class-of-service classifier ieee-802.1]** hierarchy level.

default—The default mapping.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying a Classifier to a Subscriber Interface in a Dynamic Profile on page 981
- *classifiers (Definition)*

ieee-802.1 (Dynamic Rewrite Rules)

Syntax ieee-802.1 (*rewrite-name* | default) vlan-tag (outer | outer-and-inner);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]

Release Information Statement introduced in Junos OS Release 10.1.

Description Apply an IEEE-802.1 rewrite rule to a subscriber interface in a dynamic profile.

Options *rewrite-name*—Name of a **rewrite-rules** mapping configured at the **[edit class-of-service rewrite-rules ieee-802.1]** hierarchy level.

default—The default mapping.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile on page 980
- *rewrite-rules*

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 ANCP 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 ANCP, but you want one or more neighbors to run ANCP 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 ANCP on page 1332
- Configuring ANCP Neighbors on page 1333

 Copyright © 2013, Juniper Networks, Inc.

if-exceeding (Policer)

Syntax

```
if-exceeding {
   (bandwidth-limit bps | bandwidth-percent number);
   burst-size-limit bytes;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall policer *policer-name*],
[edit firewall policer *policer-name*],
[edit logical-systems *logical-system-name* firewall policer *policer-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... policer *policer-name*]** hierarchy level introduced in Junos OS Release 11.4.

Description

Configure rate limits for a single-rate two-color policer.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Two-Color Policer Configuration Overview*
- *Hierarchical Policer Configuration Overview*
- *Basic Single-Rate Two-Color Policers*
- *Bandwidth Policers*
- *Filter-Specific Counters and Policers*
- *Prefix-Specific Counting and Policing Actions*
- *Multifield Classification*
- *Policer Overhead to Account for Rate Shaping in the Traffic Manager*
- *Hierarchical Policers*

Copyright © 2013, Juniper Networks, Inc.

if-exceeding (Hierarchical Policer)

Syntax

```
if-exceeding {
   bandwidth-limit bps;
   burst-size-limit bytes;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall hierarchical-policer aggregate],
[edit dynamic-profiles *profile-name* firewall hierarchical-policer premium],
[edit firewall hierarchical-policer aggregate],
[edit firewall hierarchical-policer premium]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... aggregate]** and **[edit dynamic-profiles ... premium]** hierarchy level introduced in Junos OS Release 11.4.

Description

For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify bandwidth and burst limits for a premium or aggregate component of a hierarchical policer.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Hierarchical Policer Configuration Overview*
- *Hierarchical Policers*
- aggregate (Hierarchical Policer) on page 1447
- bandwidth-limit (Hierarchical Policer) on page 1475
- burst-size-limit (Hierarchical Policer) on page 1484
- hierarchical-policer on page 1675
- premium (Hierarchical Policer) on page 1894

Copyright © 2013, Juniper Networks, Inc.

igmp (Dynamic Profiles)

Syntax

```
igmp {
  interface interface-name {
    accounting;
    disable;
    group-limit policy-name;
    group-policy;
    immediate-leave;
    no-accounting;
    oif-map map-name;
    passive <allow-receive> <send-general-query> <send-group-query>;
    promiscuous-mode;
    ssm-map ssm-map-name;
    static {
      group group {
        source source;
      }
    }
    version version;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* protocols],
[edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information

Statement introduced in Junos OS Release 9.2.

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).

Options

The statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For general information about configuring IGMP, see the *Multicast Protocols Feature Guide for Routing Devices*
- For information about enabling IGMP, see “Enabling IGMP” in the *Multicast Protocols Feature Guide for Routing Devices*

ignore

Syntax

```
ignore {
   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

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 RADIUS Server Parameters for Subscriber Access on page 36

immediate-leave (Dynamic IGMP Interface)

Syntax immediate-leave;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*],

Release Information Statement introduced in Junos OS Release 9.2.

Description Enable the routing device to leave the multicast group immediately after the last host leaves the multicast group.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about configuring IGMP immediate leave, see "*Specifying Immediate-Leave Host Removal for IGMP*" in the *Multicast Protocols Feature Guide for Routing Devices*

 Copyright © 2013, Juniper Networks, Inc.

immediate-leave (Dynamic MLD Interface)

Syntax immediate-leave;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description The immediate leave setting is useful for minimizing the leave latency of MLD memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows MLD to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending MLD group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the MLD leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both MLD version 1 and MLD version 2.

NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Specifying Immediate-Leave Host Removal for MLD*

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 36
- Configuring Per-Subscriber Session Accounting on page 29

include-ipv6 (Gx-Plus)

Syntax include-ipv6;

Hierarchy Level [edit access gx-plus global]

Release Information Statement introduced in Junos OS Release 12.3.

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 551
- Configuring Gx-Plus on page 549

 Copyright © 2013, Juniper Networks, Inc.

inet (Subscriber Secure Policy)

Syntax

```
inet {
  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;
    }
  }
}
```

Hierarchy Level

[edit services radius-flow-tap policy *policy-name*]

Release Information

Statement introduced in Junos OS Release 12.3.

Description

Specify the inet family for the policy that is applied to mirrored packets sent to a mediation device.

The remaining statements are explained separately.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

inet6 (Subscriber Secure Policy)

Syntax

```
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;
    }
  }
}
```

Hierarchy Level

[edit services radius-flow-tap policy *policy-name*]

Release Information

Statement introduced in Junos OS Release 12.3.

Description

Specify the inet6 family for the policy that is applied to mirrored packets sent to a mediation device.

The remaining statements are explained separately.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

Copyright © 2013, Juniper Networks, Inc.

inet-precedence (Dynamic Classifiers)

Syntax inet-precedence (*classifier-name* | default);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers]

Release Information Statement introduced in Junos OS Release 10.1.

Description Apply a IPv4 precedence classifier to a subscriber interface in a dynamic profile.

Options *classifier-name*—Name of a **classifier** mapping configured at the **[edit class-of-service classifier ieee-802.1]** hierarchy level.

default—The default mapping.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying a Classifier to a Subscriber Interface in a Dynamic Profile on page 981
- *classifiers (Definition)*

inet-precedence (Dynamic Rewrite Rules)

Syntax inet-precedence (*rewrite-name* | default);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]

Release Information Statement introduced in Junos OS Release 10.1.

Description Apply a IPv4 precedence rewrite rule.

Options *rewrite-name*—Name of a **rewrite-rules** mapping configured at the **[edit class-of-service rewrite-rules inet-precedence]** hierarchy level.

default—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile on page 980
- *rewrite-rules*

inline-services (PIC level)

Syntax

```
inline-services {
  bandwidth (1g | 10g);
}
```

Hierarchy Level

[edit chassis fpc *slot-number* pic *number*]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Enable inline services on PICs residing on MPCs. To enable inline services that are specified at the fpc level, see configuration statement *inline-services (FPC Level)*

The remaining statement is explained separately.

Options

The option is described separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling Inline Service Interfaces on page 427
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

inner-tag-protocol-id (Dynamic VLANs)

Syntax

inner-tag-protocol-id *tpids*;

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*input-vlan-map],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*output-vlan-map]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

For dynamic VLAN interfaces, configure the IEEE 802.1Q TPID value to rewrite for the inner tag. All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id *tpids*]** hierarchy level.

Default

If the **inner-tag-protocol-id** statement is not configured, the TPID value is 0x8100.

Options

tpids—TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal format.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Inner and Outer TPIDs and VLAN IDs*

Copyright © 2013, Juniper Networks, Inc.

inner-vlan-id (Dynamic VLANs)

Syntax inner-vlan-id *number*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*input-vlan-map],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*output-vlan-map]

Release Information Statement introduced in Junos OS Release 10.4.

Description For dynamic VLAN interfaces, specify the VLAN ID to rewrite for the inner tag of the final packet.

You cannot include the **inner-vlan-id** statement with the **swap** statement, **swap-push** statement, **push-push** statement, or **push-swap** statement and the **inner-vlan-id** statement at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]** hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the **inner-vlan-id** statement you include at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Options ***number***—VLAN ID number.

Range: 0 through 4094

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring Inner and Outer TPIDs and VLAN IDs*

input (Dynamic Service Sets)

Syntax

```
input {
  service-set service-set-name {
    service-filter filter-name;
  }
  post-service-filter filter-name;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* service],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos–interface–unit" family *family* service]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* service]** hierarchy level introduced in Junos OS Release 10.1.

Description

Define the input service sets and filters to be applied to traffic by a dynamic profile. Only the Internet Protocol version 4 (IPv4) protocol family is currently supported for dynamic PPPoE logical interfaces.

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

- Dynamic Service Sets Overview on page 1151
- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197

Copyright © 2013, Juniper Networks, Inc.

input-vlan-map (Dynamic Interfaces)

Syntax

```
input-vlan-map {
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  (push | swap);
  tag-protocol-id tpid;
  vlan-id number;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

For dynamic interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.

The 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

- *Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution*

instance-role

Syntax instance-role (access | nni);

Hierarchy Level [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 11.2.

Description Define the role of the routing instance in a Layer 2 Wholesale network.

Options

access—Defines the connectivity role of the routing instance in a Layer 2 Wholesale network as an access routing instance. When defined for this role, the same process occurs as in a Layer 3 Wholesale network—when the first packet is received from a given client, authentication for the client initiates with an external entity (for example, RADIUS). If authentication is successful, a logical interface is created with the appropriate outer and inner VLAN tags for that client.

nni—Defines the connectivity role of the routing instance in a Layer 2 Wholesale network as a network to network interface (NNI) routing instance. When defined for this role, only outer VLAN tags are learned. In addition, when the NNI routing instance receives a response from the ISP, the packets are forwarded to the appropriate client, provided the packet has the same two tags that were verified during authentication.

NOTE: **If you connect an access node or MSAN device to a router participating in the Layer 2 Wholesale network in an NNI role, you must create a new routing instance of type l2backhaul-vpn with an instance role of type access for that connection.**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers*
- *Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers*
- *Junos OS Subscriber Management and Services Library*

Copyright © 2013, Juniper Networks, Inc.

interface (DHCP Local Server)

Syntax

```
interface interface-name {
  exclude;
  overrides {
    client-discover-match <option60-and-option82>;
    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 only supports static DHCP configurations. For additional information about how to configure IRB, see the *Ethernet Networking Feature Guide for MX Series Routers*.

Options

exclude—Exclude an interface or a range of interfaces from the group. This option and the **overrides** option are mutually exclusive.

interface-name—Name of the interface. You can repeat this option multiple times.

upto-interface-name—Upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the ***upto-interface-name*** must be the same as the device name of the ***interface-name***.

The remaining statements are explained separately.

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 200
- Grouping Interfaces with Common DHCP Configurations on page 217
- Using External AAA Authentication Services with DHCP on page 214

Copyright © 2013, Juniper Networks, Inc.

interface (DHCP Relay Agent)

Syntax

```
interface interface-name {
  exclude;
  overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match <option60-and-option82>;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    proxy-mode;
    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.

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. For additional information about how to configure IRB, see the *Ethernet Networking Feature Guide for MX Series Routers*.

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.

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 276
- Grouping Interfaces with Common DHCP Configurations on page 217
- Using External AAA Authentication Services with DHCP on page 214

Copyright © 2013, Juniper Networks, Inc.

interface (Dynamic IGMP)

Syntax

```
interface interface-name {
  accounting;
  disable;
  group-policy;
  immediate-leave
  no-accounting;
  oif-map;
  passive;
  promiscuous-mode;
  ssm-map ssm-map-name;
  static {
    group group {
      source source;
    }
  }
  version version;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* protocols igmp]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Enable IGMP on an interface and configure interface-specific properties.

Options

interface-name—Variable for the interface. Specify the interface variable ($junos-interface-name) to indicate that the dynamic profile chooses an interface for the accessing DHCP client.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about configuring IGMP interfaces, see “Enabling IGMP” in the *Multicast Protocols Feature Guide for Routing Devices*

interface (Dynamic Interface Sets)

Syntax

```
interface interface-name {
  unit logical unit number {
    advisory-options {
      downstream-rate rate;
      upstream-rate rate;
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces interface-set *interface-set-name*]

Release Information

Statement introduced in Junos OS Release 12.2.

Description

Add a subscriber interface to a dynamic interface set.

In a dynamic profile that defines an agent circuit identifier (ACI) interface set, observe the following guidelines when you use the **interface** statement:

- Use the predefined dynamic interface variable **$junos-interface-ifd-name** to represent the interface name. Do not use a specific interface name, such as **demux0**, when defining an ACI interface set.
- Do not include the **unit *logical-unit-number*** statement.

Options

interface-name—Either the specific name of the interface to include in the interface set, or the predefined dynamic interface variable **$junos-interface-ifd-name**. The interface variable is dynamically replaced with the interface that the DHCP or PPPoE subscriber accesses when connecting to the router.

The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Defining Agent Circuit Identifier Interface Sets on page 735
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring an Interface Set of Subscribers in a Dynamic Profile on page 1028
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704

Copyright © 2013, Juniper Networks, Inc.

interface (Dynamic MLD)

Syntax

```
interface interface-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;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* protocols mld]

Release Information

Statement introduced in Junos OS Release 10.1.

Description

Enable MLD on a dynamic interface and configure interface-specific properties.

Options

interface-name—Variable for the interface. Specify the interface variable ($junos-interface-name) to indicate that the dynamic profile chooses an interface for the accessing client.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Enabling MLD*

interface (Dynamic Routing Instances)

Syntax interface *interface-name*;

Hierarchy Level [edit dynamic-profiles *profile-name*routing-instances *routing-instance-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Assign the specified interface to the dynamically created routing instance.

Options ***interface-name***—The interface name variable (*$junos-interface-name*). The interface name variable is dynamically replaced with the interface the accessing client uses when connecting to the router.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *[edit routing-instances] Hierarchy Level*
- *Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution*

 Copyright © 2013, Juniper Networks, Inc.

interface (Dynamic Router Advertisement)

Syntax

```
interface interface-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);
  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.

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*

interface (Dynamic Routing Options)

Syntax

```
interface interface-names {
  no-qos-adjust;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-options multicast],
[edit dynamic-profiles *profile-name* routing-instances *routing-instance-name* routing-options multicast]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Define the maximum bandwidth for a dynamic interface on which you want to apply bandwidth management.

Options

interface-name—Names of the physical or logical interface. For details about specifying interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 690
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 692

interface (L2TP Service Interfaces)

Syntax

interface *service-interface-name*;

Hierarchy Level

[edit services service-device-pools pool *pool-name*]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Specify a service interface assigned to a service interface pool. You specify more than one interface for each pool; the interfaces are used by an L2TP tunnel group to balance traffic loads.

Options

service-interface-name—Name of the service interface.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 430
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429

Copyright © 2013, 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 500
- Creating a Static Subscriber Group on page 504

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 500
- Configuring the Static Subscriber Global Username on page 503
- Configuring the Static Subscriber Group Username on page 507

Copyright © 2013, 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

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 221
- Overriding Default DHCP Local Server Configuration Settings on page 220

Copyright © 2013, 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

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 1548
- Extended DHCP Relay Agent Overview on page 276
- Group-Specific DHCP Relay Options on page 292
- Overriding the Default DHCP Relay Configuration Settings on page 293

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 232

Copyright © 2013, 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 47
- RADIUS Server Options for Subscriber Access on page 41

Copyright © 2013, Juniper Networks, Inc.

interface-id

Syntax interface-id *interface-id*;

Hierarchy Level [edit access group-profile *profile-name* l2tp],
[edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ike],
[edit access profile *profile-name* client *client-name* l2tp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the interface identifier.

Options ***interface-id***—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the **[edit interfaces *interface-name* unit *local-unit-number* dial-options]** hierarchy level. For more information about the interface ID, see the *Junos OS Services Interfaces Library for Routing Devices*.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring L2TP Properties for a Client-Specific Profile*
- *Configuring PPP Properties for a Client-Specific Profile*
- *Configuring an IKE Access Profile*
- Configuring an L2TP Access Profile on the LNS on page 422

 Copyright © 2013, Juniper Networks, Inc.

interface-name (DHCP Local Server)

Syntax

interface-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 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 239

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 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.

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 239

Copyright © 2013, 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.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring ANCP on page 1332
- Associating an Access Node with Subscribers for ANCP Operations on page 1334

interface-set (Dynamic CoS)

Syntax

```
interface-set interface-set-name {
  interface interface-name {
    unit logical-unit-number;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

For MX Series routers with enhanced queuing DPCs or MPC/MIC modules, configure an interface set for dynamic CoS.

Options

interface-set *interface-set-name*—Name of the scheduler to be configured or one of the following Junos OS predefined variables:

- $junos-interface-set-name—Predefined variable that, when used, is replaced with the interface-set obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.
- $junos-svlan-interface-set-name—Locally generated interface set name for use by dual-tagged VLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is ***physical_interface_name - outer_VLAN_tag***.
- $junos-tagged-vlan-interface-set-name—Locally generated interface set name used for grouping logical interfaces stacked over logical stacked VLAN demux interfaces for either a 1:1 (dual-tagged; individual client) VLAN or N:1 (single tagged; service) VLAN. The format of the generated variable differs with VLAN type. For dual-tagged (client) VLANs, the format of the generated variable is ***physical_interface_name - outer_VLAN_tag - inner_VLAN_tag***. For single tagged (service) VLAN, the format of the generated variable is ***physical_interface_name - VLAN_tag***.

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

- CoS for Interface Sets of Subscribers Overview on page 1019
- Configuring an Interface Set of Subscribers in a Dynamic Profile on page 1028
- Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile on page 1065

Copyright © 2013, Juniper Networks, Inc.

interface-set (Dynamic VLAN Interface Sets Association)

Syntax

```
interface-set interface-set-name {
  interface interface-name {
    unit logical-unit-number {
      advisory-options {
        downstream-rate rate;
        upstream-rate rate;
      }
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces]

Release Information

Statement introduced in Junos OS Release 12.2.

Description

For MX Series routers with MPC/MIC modules that face the access side of the network, associate an agent circuit identifier (ACI) interface set with a dynamic VLAN subscriber interface for DHCP or PPPoE subscribers. To associate an ACI interface set with a dynamic subscriber interface, you must include the **interface-set** stanza in the dynamic profile that defines the logical subscriber interface.

An ACI interface set is a logical collection of subscriber interfaces that originate at the same household or on the same access-loop port.

Options

- *interface-set-name*—Name of the ACI interface set, which is represented in a dynamic profile for a subscriber interface by the Junos OS predefined variable **$junos-interface-set-name**.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic VLAN Subscriber Interfaces Based on Agent Circuit Identifier Information on page 740
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704

interface-set (Dynamic VLAN Interface Sets Definition)

Syntax

```
interface-set interface-set-name {
  interface interface-name;
  pppoe-underlying-options {
    max-sessions number;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces]

Release Information

Statement introduced in Junos OS Release 12.2.

Description

For MX Series routers with MPC/MIC modules that face the access side of the network, configure an agent circuit identifier (ACI) interface set for the creation of dynamic VLAN subscriber interfaces for DHCP or PPPoE subscribers based on ACI information. An ACI interface set is a logical collection of subscriber interfaces that originate at the same household or on the same access-loop port.

To configure an ACI interface set for dynamic VLAN subscriber interfaces, you must include the **interface-set** stanza in the dynamic profile that defines the ACI interface set.

Options

- *interface-set-name*—Name of the ACI interface set, which is represented in a dynamic profile by the Junos OS predefined variable **$junos-interface-set-name**.

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

- Defining Agent Circuit Identifier Interface Sets on page 735
- Clearing Agent Circuit Identifier Interface Sets on page 748
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704

Copyright © 2013, Juniper Networks, Inc.

interface-shared

Syntax interface-shared;

Hierarchy Level [edit firewall family *family-name* filter *filter-name*],
[edit dynamic-profiles *profile-name*firewall family *family-name* filter *filter-name*]

Release Information Statement introduced in Junos OS Release 12.2.

Description Set the interface-shared attribute for a firewall filter.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Firewall Filters Overview on page 1132
- Classic Filters Overview on page 1133
- Basic Classic Filter Syntax on page 1135

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.
- **notice**—Match notice messages about conditions requiring special handling.

Copyright © 2013, Juniper Networks, Inc.

- **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.

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 260

interfaces (Subscriber Secure Policy)

Syntax

interfaces *interface-name*;

Hierarchy Level

[edit services radius-flow-tap]

Release Information

Statement introduced in Junos OS Release 9.4.

Description

Specify tunnel interfaces that are used to send mirrored packets to a mediation device.

Options

interface-name—Name of the interface.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

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 ANCP.

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.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring ANCP on page 1332
- Associating an Access Node with Subscribers for ANCP Operations on page 1334

Copyright © 2013, Juniper Networks, Inc.

interfaces (Dynamic CoS Definition)

Syntax

```
interfaces {
  interface-name {
    unit logical-unit-number {
      classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        ieee-802.1 (classifier-name | default) vlan-tag (inner | outer)
        inet-precedence (classifier-name | default);
      }
      output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
      rewrite-rules {
        dscp (rewrite-name | default);
        dscp-ipv6 (rewrite-name | default);
        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
        inet-precedence (rewrite-name | default);
      }
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Configure interface-specific CoS properties for incoming packets.

Options

interface-name—Either the specific name of the interface you want to assign to the dynamic profile or the interface variable ($junos-interface-ifd-name). The interface variable is dynamically replaced with the interface the client accesses when connecting to the router.

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

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile on page 979

Copyright © 2013, Juniper Networks, Inc.

interfaces (Static and Dynamic Subscribers)

Syntax

```
interfaces {
  interface-name {
    unit logical-unit-number {
      auto-configure {
        agent-circuit-identifier {
          dynamic-profile profile-name;
        }
      }
      family family {
        access-concentrator name;
        address address;
        duplicate-protection;
        dynamic-profile profile-name;
        filter {
          adf {
            counter;
            input-precedence precedence;
            not-mandatory;
            output-precedence precedence;
            rule rule-value;
          }
          input filter-name (
            precedence precedence;
            shared-name filter-shared-name;
          }
          output filter-name {
            precedence precedence;shared-name filter-shared-name;
          }
        }
        max-sessions number;
        max-sessions-vsa-ignore;
        rpf-check {
          mode loose;
        }
        service {
          input {
            service-set service-set-name {
              service-filter filter-name;
            }
            post-service-filter filter-name;
          }
          output {
            service-set service-set-name {
              service-filter filter-name;
            }
          }
        }
        service-name-table table-name
        short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
          maximum-seconds>;
        unnumbered-address interface-name <preferred-source-address address>;
      }
```


Copyright © 2013, Juniper Networks, Inc.

```
        filter {
          input filter-name;
          shared-name filter-shared-name;
          output filter-name;
          shared-name filter-shared-name;
        }
        ppp-options {
          chap;
          pap;
        }
        proxy-arp;
        vlan-id;
        vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
      }
      vlan-tagging;
    }
    interface-set interface-set-name {
      interface interface-name {
        unit logical unit number {
          advisory-options {
            downstream-rate rate;
            upstream-rate rate;
          }
        }
      }
      pppoe-underlying-options {
        max-sessions number;
      }
    }
    demux0 {
      unit logical-unit-number {
        demux-options {
          underlying-interface interface-name
        }
        family family {
          access-concentrator name;
          address address;
          duplicate-protection;
          dynamic-profile profile-name;
          demux-source {
            source-prefix;
          }
          filter{
            input filter-name (
              precedence precedence;
              shared-name filter-shared-name;
            }
            output filter-name {
              precedence precedence;
              shared-name filter-shared-name;
            }
          }
          mac-validate (loose | strict):
          max-sessions number;
          max-sessions-vsa-ignore;
          rpf-check {
```

```
            fail-filter filter-name;
            mode loose;
          }
          service-name-table table-name
          short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
            maximum-seconds>;
          unnumbered-address interface-name <preferred-source-address address>;
        }
        filter {
          input filter-name;
          output filter-name;
        }
        vlan-id number;
        vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
      }
    }
    pp0 {
      unit logical-unit-number {
        keepalives interval seconds;
        no-keepalives;
        pppoe-options {
          underlying-interface interface-name;
          server;
        }
        ppp-options {
          authentication [ authentication-protocols ];
          chap {
            challenge-length minimum minimum-length maximum maximum-length;
          }
          pap;
        }
        family inet {
          unnumbered-address interface-name;
          address address;
          service {
            input {
              service-set service-set-name {
                service-filter filter-name;
              }
              post-service-filter filter-name;
            }
            output {
              service-set service-set-name {
                service-filter filter-name;
              }
            }
          }
          filter {
            input filter-name {
              precedence precedence;
              shared-name filter-shared-name;
            }
            output filter-name {
              precedence precedence;
              shared-name filter-shared-name;
            }
```


Copyright © 2013, Juniper Networks, Inc.

```
            }
          }
        }
      }
    }
```

Hierarchy Level [edit dynamic-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Define interfaces for dynamic profiles.

Options ***interface-name***—The interface variable (**$junos-interface-ifd-name**). The interface variable is dynamically replaced with the interface the DHCP client accesses when connecting to the router.

NOTE: **Though we do not recommend it, you can also enter the specific name of the interface you want to assign to the dynamic profile.**

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Static Subscriber Interfaces in Dynamic Profiles on page 765
- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles on page 899
- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Subscriber Interface Overview on page 757
- Relationship Between Subscribers and Interfaces in an Access Network on page 5
- For general information about configuring static interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*
- For information about static IP demux interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

interface-specific (Dynamic Firewalls)

Syntax interface-specific;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall family *family* fast-update-filter *filter-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure interface-specific names for firewall counters that are based on fast update filters.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Fast Update Filters on page 1183

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 168
- Configuring Static Address Assignment on page 171

 Copyright © 2013, Juniper Networks, Inc.

ip-address-change-notify

Syntax ip-address-change-notify *message*;

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the Unisphere-IPv4-release-control VSA in RADIUS messages. When enabled, the BNG includes Unisphere-Ipv4-release-control VSA in the Access-Request that is sent during on-demand IP address allocation and in the immediate Interim-Accounting messages that are sent to report an address change. Disabled by default, there is no effect when on-demand IP address allocation or deallocation is not configured. An change takes effect immediately. It is optional to specify the message, but if specified, the message is inserted into Unisphere-Ipv4-release-control VSA. Otherwise, a default value (NO MESSAGE) is be inserted into the VSA.

Options ***message***—VSA message.

Range: 1 through 32 characters,

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Saving IPv4 Addresses for Dual-Stack PPP Subscribers*

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 215
- Extended DHCP Local Server Overview on page 200
- Address-Assignment Pools Overview on page 167
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

Copyright © 2013, Juniper Networks, Inc.

ip-reassembly (L2TP)

Syntax

```
ip-reassembly {
   service-set service-set-name;
}
```

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 13.1.

Description Associate the reassembly service-set with the L2TP service.

NOTE: **The service set must be defined at the [edit services] hierarchy level.**

Options **service-set *service-set-name***—Identifies the service set to be associated with the L2TP service.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- IP Packet Fragment Reassembly for L2TP Overview on page 390
- Configuring IP Inline Reassembly for L2TP on page 421

ip-reassembly

Syntax

```
ip-reassembly {
   profile profile-name
   rule rule-name{
      match-direction direction
   };
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the IP reassembly parameters to be applied to the L2TP server.

NOTE: Inline IP reassembly configuration does not require you to configure the profile statement. The profile configuration is used when IP reassembly is configured on services PICs.

Options **profile *profile-name***—Name of the IP reassembly profile.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring IP Inline Reassembly for L2TP on page 421
- IP Packet Fragment Reassembly for L2TP Overview on page 390

 Copyright © 2013, Juniper Networks, Inc.

ip-reassembly-rules (Service Set)

Syntax ip-reassembly-rules *rule-name*;

Hierarchy Level [edit services service-set*service-set-name*]

Release Information Statement introduced in Junos OS Release 13.1.

Description Specify one or more previously configured IP reassembly rules to associate with the service set.

NOTE: **The IP reassembly rule must be defined at the [edit services ip-reassembly rule] hierarchy level.**

Options ***rule-name***—Name of an IP reassembly rule.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring IP Inline Reassembly for L2TP on page 421
- IP Packet Fragment Reassembly for L2TP Overview on page 390

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.

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 491

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 491
- Configuring the JSRC Partition on page 492

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 1341
- Configuring ANCP on page 1332

 Copyright © 2013, Juniper Networks, Inc.

keepalive

Syntax keepalive *seconds*;

Hierarchy Level [edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the keepalive interval for an L2TP tunnel.

Options ***seconds***—Time period that must elapse before the Junos OS checks the status of the Point-to-Point Protocol (PPP) session by sending an echo request to the peer.

For L2TP on MX Series routers, the minimum recommended interval is 30 seconds. A value of 0 disables generation of keepalive messages from the LNS.

Range: 0 through 32,767 seconds

Default: 30 seconds

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring PPP Properties for a Client-Specific Profile*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 418

keepalives

Syntax keepalives <interval *seconds*> <down-count *number*> <up-count *number*>;

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation.

For ATM2 IQ interfaces only, you can enable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

Default Sending of keepalives is enabled by default. The default keepalive interval is 10 seconds for PPP, Frame Relay, or Cisco HDLC. The default down-count is 3 and the default up-count is 1 for PPP or Cisco HDLC.

Options **down-count *number***—The number of keepalive packets a destination must fail to receive before the network takes down a link.

Range: 1 through 255

Default: 3

interval *seconds*—The time in seconds between successive keepalive requests.

Range: 1 through 32767 seconds

Default: 10 seconds

up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up.

Range: 1 through 255

Default: 1

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Keepalives*
- *Configuring Frame Relay Keepalives*
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 419

 Copyright © 2013, Juniper Networks, Inc.

keepalives (Dynamic Profiles)

Syntax

```
keepalives {
   interval seconds;
   }
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit *logical-unit-number*]
[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit”]
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit”]** hierarchy level introduced in Junos OS Release 10.1.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]** hierarchy level introduced in Junos OS Release 12.2.

Description

Specify the keepalive interval in a PPP dynamic profile.

Default

Sending of keepalives is enabled by default.

Options

interval *seconds*—The time in seconds between successive keepalive requests.

Range: 1 through 32767 seconds

Default: 30 seconds for LNS-based PPP sessions. 10 seconds for all other PPP sessions.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring Dynamic Authentication for PPP Subscribers on page 365
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 419

Copyright © 2013, Juniper Networks, Inc.

key

Syntax

key (hex | ascii) *string*;

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* services mobile-ip peer nai *name@domain* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer nai *name@domain* spi *hexadecimal-value*],
[edit routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit routing-instances *routing-instances-name* services mobile-ip peer nai *name@domain* spi *hexadecimal-value*],
[edit services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit services mobile-ip peer nai *name@domain* spi *hexadecimal-value*]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Configure the authentication key for the security association, in either HEX or ASCII format. The resulting 128–bit key is specified as a hexadecimal number with each character in the range 0x0–0xF.

Options

hex *string*—Key specified in HEX format

ascii *string*—Key specified in ASCII format

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

Copyright © 2013, Juniper Networks, Inc.

l2tp

Syntax

```
l2tp {
  destination
    lockout-timeout seconds;
  }
  destruct-timeout seconds;
  disable-calling-number-avp;
  disable-failover-protocol;
  enable-snmp-tunnel-statistics;
  fail-over-within-preference;
  ip-reassembly;
  rx-connect-speed-when-equal;
  traceoptions {
    debug-level level;
    file filename <files number> <match regular-expression > <size maximum-file-size>
      <world-readable | no-world-readable>;
    filter {
      protocol name;
      user-name username;
    }
    flag flag;
    interfaces interface-name {
      debug-level severity;
      flag flag;
    }
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
  tunnel {
    assignment-id-format (assignment-id |client-server-id);
    idle-timeout seconds;
    retransmission-count-established count;
    retransmission-count-not-established count;
    tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port);
    rx-window-size packets;
  }
  tunnel-group group-name {
    aaa-access-profile profile-name;
    dynamic-profile profile-name;
    hello-interval seconds;
    hide-avps;
    l2tp-access-profile profile-name;
    local-gateway address address;
    maximum-send-window packets;
    ppp-access-profile profile-name;
    receive-window packets;
    retransmit-interval seconds;
    service-device-pool pool-name;
    service-interface interface-name;
    syslog {
      host hostname {
        facility-override facility-name;
        log-prefix prefix-value;
```

```
            services severity-level;
          }
        }
        tos-reflect;
        tunnel-switch-profile profile-name;
        tunnel-timeout seconds;
      }
      tunnel-switch-profile profile-name;
      tx-connect-speed-method method;
      weighted-load-balancing;
    }
```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4.
Support for LAC on MX Series routers introduced in Junos OS Release 10.4.
Support for LNS on MX Series routers introduced in Junos OS Release 11.4.

Description Configure L2TP services to establish PPP tunnels across a network.

The remaining statements are explained separately.

NOTE: **Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.**

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Layer 2 Tunneling Protocol Overview*
- L2TP for Subscriber Access Overview on page 379

 Copyright © 2013, Juniper Networks, Inc.

l2tp-access-profile

Syntax l2tp-access-profile *profile-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the profile used to validate all L2TP connection requests to the local gateway address.

Options ***profile-name***—Identifier for the L2TP connection profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Access Profiles for L2TP Tunnel Groups*
- Configuring an L2TP Access Profile on the LNS on page 422

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 276
- dhcp-relay on page 1548

Copyright © 2013, Juniper Networks, Inc.

lcp-renegotiation

Syntax lcp-renegotiation;

Hierarchy Level [edit access group-profile *profile-name* l2tp],
[edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the L2TP network server (LNS) so it renegotiates the link control protocol (LCP) with the PPP client. When LCP renegotiation is disabled, LNS uses the pre-negotiated LCP parameters between the L2TP access concentrator (LAC) and PPP client to set up the session. When LCP renegotiation is enabled, authentication is also renegotiated.

NOTE: **This statement is not supported at the [edit access group-profile l2tp] hierarchy level for L2TP LNS on MX Series routers.**

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring L2TP Properties for a Client-Specific Profile*
- Configuring an L2TP Access Profile on the LNS on page 422

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.

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 the secondary address-assignment pool that is linked to a primary address-assignment pool. The secondary pool provides backup pool for local address assignment.

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 167
- Configuring Address-Assignment Pools on page 168
- Configuring Address-Assignment Pool Linking on page 170

local-gateway address

Syntax local-gateway address *address*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the local (LNS) IP address for L2TP tunnel.

Options ***address***—Local IP address; corresponds to the IP address that is used by LACs to identify the LNS. When the LAC is an MX Series router, this address matches the remote gateway address configured in the LAC tunnel profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Local Gateway Address and PIC*.
- *Configuring L2TP Tunnel Groups*
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429

 Copyright © 2013, Juniper Networks, Inc.

liveness-detection

Syntax

```
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode(automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
```

Hierarchy Level

[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit system services dhcp-local-server group *group-name*],
[edit system services dhcp-local-server dhcpv6 group *group-name*],
[edit forwarding-options dhcp-relay group *group-name*],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name*]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure bidirectional failure detection timers and authentication criteria for static routes.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- DHCP Liveness Detection Overview on page 250
- Configuring Detection of DHCP Local Server Client Connectivity on page 251
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 318
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269

Copyright © 2013, Juniper Networks, Inc.

- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

lockout-timeout (L2TP Destination Lockout)

Syntax lockout-timeout *seconds*;

Hierarchy Level [edit services l2tp destination]

Release Information Statement introduced in Junos OS Release 13.2.

Description Set the duration of the timeout period for which all future destinations are locked out, meaning that they are not considered for selection when a new tunnel is created. This statement does not affect destinations that are currently locked out.

Options ***seconds***—Length of the period during which the destination is locked out.

Range: 60 through 3600 seconds

Default: 300 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the L2TP Destination Lockout Timeout on page 414
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

 Copyright © 2013, 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 323

logical-bandwidth-policer

Syntax logical-bandwidth-policer;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall policer *policer-name*],
[edit firewall policer *policer-name*],
[edit logical-systems *logical-system-name* firewall policer *policer-name*]

Release Information Statement introduced in Junos OS Release 8.2.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... policer *policer-name*]** hierarchy level introduced in Junos OS Release 11.4.

Description For a policer with a bandwidth limit configured as a percentage (using the bandwidth-percent statement), specify that the percentage be based on the shaping rate defined on the logical interface, rather than on the media rate of the physical interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Bandwidth Policers*
- *Configuring Logical Bandwidth Policers*
- bandwidth-percent on page 1478 statement
- **interface-specific** statement

logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces)

Syntax logical-interface-fpc-redundancy;

Hierarchy Level [edit interfaces ae*number* aggregated-ether-options]

Release Information Statement introduced in Junos OS Release 11.2.

Description Provide module redundancy for demux subscribers on aggregated Ethernet bundles configured with targeted distribution. Backup links for a subscriber are chosen on a different EQ DPC or MPC from the primary link, based on the link with the fewest number of subscribers among the links on different modules. If all links are on a single module when this is configured, backup links are not provisioned.

By default, link redundancy is provided for the aggregated Ethernet bundle.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface on page 829
- *Configuring Module Redundancy for a Virtual Chassis*

 Copyright © 2013, Juniper Networks, Inc.

logical-interface-policer

Syntax logical-interface-policer;

Hierarchy Level
[edit dynamic-profiles *profile-name* firewall policer *policer-name*],
[edit dynamic-profiles *profile-name* firewall three-color-policer *name*],
[edit firewall atm-policer*atm-policer-name*]
[edit firewall policer *policer-name*],
[edit firewall policer *policer-template-name*],
[edit firewall three-color-policer *policer-name*],
[edit logical-systems *logical-system-name* firewall policer *policer-name*],
[edit logical-systems *logical-system-name* firewall three-color-policer *name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Support at the **[edit firewall three-color-policer *policer-name*]** hierarchy level introduced in Junos OS Release 8.2.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... policer *policer-name*] and [edit dynamic-profiles ... three-color-policer *name*]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure a logical interface policer.

NOTE: Starting in Junos OS Release 12.2R2, on T Series Core Routers only, you can configure an MPLS LSP policer for a specific LSP to be shared across different protocol family types. You must include the logical-interface-policer statement to do so.

Required Privilege Level
firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Two-Color and Three-Color Logical Interface Policers*
- *Traffic Policer Types*
- *Configuring Tricolor Marking Policers*
- action on page 1432
- *Configuring Gigabit Ethernet Two-Color and Tricolor Policers*
- *action*

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 471
- Configuring Diameter Peers on page 473

 Copyright © 2013, 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 471
- Configuring the Diameter Transport on page 474

logical-system (Tunnel Profile)

Syntax logical-system *logical-system-name*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify a logical system for a tunnel. When you specify a logical system, you must also specify a routing instance.

Options ***logical-system-name***— Name of the logical system.

Default: Logical system *default*

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

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 214

Copyright © 2013, 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 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.

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.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214
- Creating Unique Usernames for DHCP Clients on page 239

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 500
- Configuring the Static Subscriber Global Username on page 503
- Configuring the Static Subscriber Group Username on page 507

Copyright © 2013, Juniper Networks, Inc.

loss-priority (Dynamic Schedulers)

Syntax loss-priority (any | low | medium-low | medium-high | high);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name* drop-profile-map]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify a loss priority to which to apply a drop profile in a dynamic profile. The drop profile map sets the drop profile for a specific PLP and protocol type. The inputs for the map are the PLP designation and the protocol type. The output is the drop profile.

Options

any—The drop profile applies to packets with any PLP.

high—The drop profile applies to packets with high PLP.

medium-high—The drop profile applies to packets with medium-high PLP.

medium-low—The drop profile applies to packets with medium-low PLP.

low—The drop profile applies to packets with low PLP.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973

loss-priority high then discard (Three-Color Policer)

Syntax

loss-priority high then discard;

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall three-color-policer *name* action],
[edit firewall three-color-policer *policer-name* action],
[edit logical-systems *logical-system-name* firewall three-color-policer *policer-name* action]

Release Information

Statement introduced before Junos OS Release 8.2.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... action]** hierarchy level introduced in Junos OS Release 11.4.

Description

For packets with high loss priority, discard the packets. The loss priority setting is implicit and is not configurable. Include this statement if you do not want the local router to forward packets that have high packet loss priority.

For single-rate three-color policers, the Junos OS assigns high loss priority to packets that exceed the committed information rate and the excess burst size.

For two-rate three-color policers, the Junos OS assigns high loss priority to packets that exceed the peak information rate and the peak burst size.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Three-Color Policer Configuration Overview*
- *Basic Single-Rate Three-Color Policers*
- *Basic Two-Rate Three-Color Policers*
- *Two-Color and Three-Color Logical Interface Policers*
- *Two-Color and Three-Color Physical Interface Policers*
- *Two-Color and Three-Color Policers at Layer 2*
- action on page 1432

Copyright © 2013, 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 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 MAC address 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 214

mac-address (DHCP Relay Agent)

Syntax

mac-address;

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 MAC address from the client PDU be 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 214

Copyright © 2013, Juniper Networks, Inc.

mac-address (Dynamic Access-Internal Routes)

Syntax mac-address *address*;

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access-internal route *subscriber-ip-address* qualified-next-hop *underlying-interface*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access-internal route *subscriber-ip-address* qualified-next-hop *underlying-interface*],
[edit dynamic-profiles routing-options access-internal route *subscriber-ip-address* qualified-next-hop *underlying-interface*]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options route *subscriber-ip-address* qualified-next-hop *underlying-interface*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* route *subscriber-ip-address* qualified-next-hop *underlying-interface*]** hierarchy levels introduced in Junos OS Release 10.1.

Description Dynamically configure the MAC address variable for an access-internal route for unnumbered interfaces such as DHCP subscriber interfaces.

Options *address*—Either the specific MAC address you want to assign to the access-internal route or the MAC address variable ($junos-subscriber-mac-address). The MAC address variable is dynamically replaced with the value supplied by DHCP when a subscriber logs in.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 692

mac-validate (Dynamic IP Demux Interface)

Syntax mac-validate (loose | strict);

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Enable IP and MAC address validation for dynamic IP demux interfaces in a dynamic profile. Supported on MX Series routers only.

Options

loose—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not match the MAC address of the tuple. Continues to forward incoming packets when the source address of the incoming packet does not match any of the trusted IP addresses.

strict—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring MAC Address Validation for Subscriber Interfaces on page 773

 Copyright © 2013, 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, subscribers when the specified type of event occurs (such as when an interface is deleted).

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.

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 232

managed-configuration (Dynamic Router Advertisement)

Syntax (managed-configuration | no-managed-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 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*

mandatory

Syntax mandatory;

Hierarchy Level [edit dynamic-profiles *profile-name* variables *variable-name*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure RADIUS to return a value for a user-defined variable. If RADIUS does not return a value for the variable, the dynamic profile fails.

When a dynamic profile has mandatory and non-mandatory variables, configure mandatory variables first in the profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring User-Defined CoS Variables in a Dynamic Service Profile on page 1000

 Copyright © 2013, 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;
  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.

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@xyz.com**, the domain map name is **xyz.com**.

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

- Configuring a Domain Map on page 181

match-direction

Syntax match-direction (input | output | input-output);

Hierarchy Level [edit services (captive-portal-content-delivery) captive-portal-content-delivery rule *rule-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the direction in which the rule match is applied.

Options **input**—Apply the rule match on the input side of the interface.

output—Apply the rule match on the output side of the interface.

input-output—Apply the rule match bidirectionally.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Redirecting HTTP Requests Overview on page 1225

match-direction (IP Reassembly Rule)

Syntax match-direction *direction*

Hierarchy Level [edit services ip-reassemblyrule*rule-name*]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the direction in which the IP reassembly rule matching is applied. The match direction is used with respect to the traffic flow through the inline services interface. You must configure a match direction for an IP reassembly rule.

Options ***direction***—Match direction. For inline IP reassembly, **input** is the only match direction supported.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring IP Inline Reassembly for L2TP on page 421
- IP Packet Fragment Reassembly for L2TP Overview on page 390

 Copyright © 2013, Juniper Networks, Inc.

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 190

match-order (Dynamic Firewalls)

Syntax match-order [*match-order*];

Hierarchy Level [edit dynamic-profiles *profile-name* firewall family *family* fast-update-filter *filter-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the match conditions and the order in which the conditions are examined. Enclose a string of multiple conditions in brackets. The router examines only the conditions you specify, and examines them in the specified order.

Options ***match-order***—One or more of the following conditions. "Fast Update Filter Match Conditions" on page 1186 describes the match conditions.

- destination-address
- destination-port
- dscp (IPv4 only)
- protocol (IPv4 only)
- source-address
- source-port

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Fast Update Filters on page 1183
- Configuring the Match Order for Fast Update Filters on page 1184
- Fast Update Filter Match Conditions on page 1186

max-advertisement-interval (Dynamic Router Advertisement)

Syntax max-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 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*

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.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

 Copyright © 2013, 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 neighbors or from a particular 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 ANCP on page 1332
- Configuring ANCP Neighbors on page 1333

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 ANCP 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 ANCP on page 1332
- Specifying How Long Processes Wait for ANCP Restart to Complete on page 1337

maximum-lease-time

Syntax maximum-lease-time *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes]

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.

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 168

maximum-sessions-per-tunnel

Syntax maximum-sessions-per-tunnel *number*;

Hierarchy Level [edit access group-profile l2tp],
[edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the maximum sessions for a Layer 2 tunnel.

NOTE: **This statement is not supported at the [edit access group-profile l2tp] hierarchy level for L2TP LNS on MX Series routers.**

Options ***number***—Maximum number of sessions for a Layer 2 tunnel.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring L2TP Properties for a Client-Specific Profile*
- Configuring an L2TP Access Profile on the LNS on page 422

 Copyright © 2013, 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 (M120, M320, MX Series routers) 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 23
- Configuring RADIUS Server Options for Subscriber Access on page 47
- *show network-access aaa statistics*
- *clear network-access aaa statistics*

max-outstanding-requests (Gx-Plus)

Syntax max-outstanding-requests *number*;

Hierarchy Level [edit access gx-plus global]

Release Information Statement introduced in Junos OS Release 11.2.

Description Limit the number of outstanding requests to the PCRF that Gx-Plus can retry when the requests are improperly answered. Too many requests risks overloading the PCRF and increases the chance of losing messages.

Options ***number***—Number of outstanding requests from Gx-Plus to the PCRF 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 551
- Configuring Gx-Plus on page 549

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 38

 Copyright © 2013, Juniper Networks, Inc.

max-sessions (Dynamic PPPoE)

Syntax max-sessions *number*;

Hierarchy Level
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces interface-set *interface-set-name* pppoe-underlying-options]
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information
Statement introduced in Junos OS Release 10.1.
Support for the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.
Support at the **[edit dynamic-profiles ... interfaces interface-set ... pppoe-underlying-options]** hierarchy level introduced in Junos OS Release 12.2.

Description Configure the maximum number of dynamic PPPoE logical interfaces that the router can activate on the underlying interface. The **max-sessions** value does not affect the maximum number of static PPPoE logical interfaces that can be configured on the underlying interface.

NOTE: The **[edit ... family pppoe]** hierarchies and the **[edit dynamic-profiles ... interfaces interface-set ... pppoe-underlying-options]** hierarchy level are supported only on MX Series routers with MPCs/MICs.

Options *number*—Maximum number of dynamic PPPoE logical interfaces (sessions) that the router can activate on the underlying interface. The default value is equal to the maximum number of PPPoE sessions supported on your routing platform. You can configure from 1 to the platform-specific default for your routing platform. Changing the **max-sessions** value has no effect on dynamic PPPoE logical interfaces that are already active.

For information about scaling values for PPPoE interfaces, access the *Subscriber Management Scaling Values (XLS)* spreadsheet from the Downloads box on the *Junos OS Subscriber Management* pathway page for the current release.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface on page 908
- Defining Agent Circuit Identifier Interface Sets on page 735

- PPPoE Maximum Session Limit Overview on page 890
- Guidelines for Using PPPoE Maximum Session Limit from RADIUS on page 892
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- For information about configuring dynamic PPPoE subscriber interfaces, see the *Junos OS Subscriber Management and Services Library*
- For information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

max-sessions (PPPoE Service Name Tables)

Syntax max-sessions *number*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the maximum number of active PPPoE sessions using either static or dynamic PPPoE interfaces that the router can establish with the specified named service, **empty** service, or **any** service entry in a PPPoE service name table. The router maintains a count of active PPPoE sessions for each service entry to determine when the maximum sessions limit has been reached.

The router uses the **max-sessions** value for a PPPoE service name table entry in conjunction with the **max-sessions** value configured for the PPPoE underlying interface, and with the maximum number of PPPoE sessions supported on your router. If your configuration exceeds any of these maximum session limits, the router is unable to establish the PPPoE session.

Options ***number***—Maximum number of active PPPoE sessions that the router can establish with the specified PPPoE service name table entry, in the range 1 to the platform-specific maximum PPPoE sessions supported for your router. The default value is equal to the maximum number of PPPoE sessions supported on your routing platform.

For information about scaling values for PPPoE interfaces, access the *Subscriber Management Scaling Values (XLS)* spreadsheet from the Downloads box on the *Junos OS Subscriber Management* pathway page for the current release.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name*
- *Configuring PPPoE Service Name Tables*
- PPPoE Maximum Session Limit Overview on page 890
- For information about configuring dynamic PPPoE subscriber interfaces, see the *Junos OS Subscriber Management and Services Library*
- For information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

max-sessions (Tunnel Profile)

Syntax max-sessions *number*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the maximum number of sessions allowed in the tunnel.

Options *number*—Maximum number of sessions allowed in the tunnel.
Range: 0 through 60,000
Default: 0

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 403

max-sessions-vsa-ignore (Static and Dynamic Subscribers)

Syntax max-sessions-vsa-ignore;

Hierarchy Level
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the router to ignore (clear) the value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks vendor-specific attribute (VSA) [26-143], and restore the PPPoE maximum session value on the underlying interface to the value configured in the CLI with the **max-sessions** statement. The PPPoE maximum session value specifies the maximum number of concurrent static or dynamic PPPoE logical interfaces (sessions) that the router can activate on the PPPoE underlying interface, or the maximum number of active static or dynamic PPPoE sessions that the router can establish with a particular service entry in a PPPoE service name table.

Default If you do not include the **max-sessions-vsa-ignore** statement, the maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA takes precedence over the PPPoE maximum session value configured with the **max-sessions** statement.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface on page 908
- PPPoE Maximum Session Limit Overview on page 890
- Guidelines for Using PPPoE Maximum Session Limit from RADIUS on page 892
- Juniper Networks VSAs Supported by the AAA Service Framework on page 98
- For information about configuring dynamic PPPoE subscriber interfaces, see the *Junos OS Subscriber Management and Services Library*
- For information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

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 38

medium (Tunnel Profile)

Syntax medium *type*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the medium type for the tunnel.

Default ipv4

Options *type*—Medium type for the tunnel. The only value currently available is **ipv4**.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 403

method

Syntax

```
method {
  bfd {
    version (0 | 1 | automatic);
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    detection-time {
      threshold milliseconds;
    }
    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
  }
}
```

Hierarchy Level

[edit system services dhcp-local-server liveness-detection],
[edit system services dhcp-local-server dhcpv6 liveness-detection],
[edit forwarding-options dhcp-relay liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection],
[edit system services dhcp-local-server group *group-name* liveness-detection],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the liveness detection method.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

Copyright © 2013, Juniper Networks, Inc.

metric (Dynamic Access-Internal Routes)

Syntax metric *route-cost*;

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*],
[edit dynamic-profiles *profile-name* routing-options access route *prefix*]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*]** hierarchy levels introduced in Junos OS Release 10.1.

Description Dynamically configure the cost for an access route.

Options *route-cost*—Either the specific cost you want to assign to the access route or either of the following cost variables:

- **$junos-framed-route-cost**—Cost of an IPv4 access route; the variable is dynamically replaced with the metric value (Subattribute 3) from the RADIUS Framed-Route attribute [22].
- **$junos-framed-route-ipv6-cost**—Cost of an IPv6 access route; the variable is dynamically replaced with the metric value (Subattribute 3) from the RADIUS Framed-IPv6-Route attribute [99].

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Access Routes for Subscriber Management on page 690

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 471
- Configuring Diameter Network Elements on page 473

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 190

 Copyright © 2013, Juniper Networks, Inc.

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*

minimum-interval

Syntax minimum-interval *milliseconds*;

Hierarchy Level
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd transmit-interval]

Release Information
Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimal-interval** and **minimum-receive-interval** statements.

Options *milliseconds* — Specify the minimum interval value for BFD liveliness detection.

Range: 1 through 255,000

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

 Copyright © 2013, Juniper Networks, Inc.

minimum-receive-interval

Syntax minimum-receive-interval *milliseconds*;

Hierarchy Level [edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the minimum interval at which the local routing device (or switch) must receive a reply from a neighbor with which it has established a BFD session.

Options ***milliseconds*** — Specify the minimum receive interval value.

Range: 1 through 255,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

mld (Dynamic Profiles)

Syntax

```
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;
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* protocols]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure interface-specific MLD values on dynamic interfaces.

Options The statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Enabling MLD*

Copyright © 2013, Juniper Networks, Inc.

mobile-ip

Syntax

```
mobile-ip {
  access-type {
    (generic | wimax);
  }
  authenticate {
    order (aaa | local);
  }
  dynamic-home-assignment {
    home-agent {
      nai (name@domain | @domain) {
        home-agent ip-address;
      }
    }
  }
  home-agent {
    enable-service interface-name;
    virtual-network {
      home-agent-address ip-address {
        registration-lifetime seconds;
        revocation-required;
        timestamp-tolerance seconds;
      }
    }
  }
  peer {
    (ip-address address | nai name@domain) {
      spi hexadecimal-value {
        algorithm (hmac-md5 | md5);
        entity-type (host | mobility-agent);
        key (hex | ascii) string;
        replay-method (none | timestamp seconds);
      }
    }
  }
  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 services],
[edit logical-systems *logical-system-name* services],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services],
[edit routing-instances *routing-instances-name* services]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* services]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services]**, and **[edit**

routing-instances *routing-instances-name* services] hierarchy levels introduced in Junos OS Release 9.5.

Description Configure Junos Mobile IP features.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Mobile IP Home Agent on page 570

mode (Dynamic Profiles)

Syntax mode loose;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family (inet) rpf-check]

Release Information Statement introduced in Junos OS Release 9.6.

Description Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.

Default If you do not include this statement, unicast RPF is in strict mode.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring Unicast RPF*

 Copyright © 2013, Juniper Networks, Inc.

multicast (Dynamic Routing Options)

Syntax

```
multicast {
  interface interface-name {
    no-qos-adjust;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-options],
[edit dynamic-profiles *profile-name* routing-instances *routing-instance-name* routing-options]

NOTE: You cannot apply a scope policy to a specific routing instance. That is, all scoping policies are applied to all routing instances. However, the scope statement does apply individually to a specific routing instance.

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Dynamically configure interface-specific multicast routing options properties.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring the Multicast Forwarding Cache*
- *Example: Configuring a Multicast Flow Map*
- *Example: Configuring Source-Specific Multicast Groups with Any-Source Override*

Copyright © 2013, Juniper Networks, Inc.

multicast-interception (Subscriber Secure Policy)

Syntax multicast-interception;

Hierarchy Level [edit services radius-flow-tap]

Release Information Statement introduced in Junos OS Release 11.4.

Description Enables subscriber secure policy to mirror IPv4 multicast traffic sent to subscribers. It enables the mirroring of multicast traffic for all subscribers on the chassis.

Mirroring of multicast traffic is supported only for subscribers in the default logical system.

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 on page 1245
- Subscriber Secure Policy Support for IPv4 Multicast Traffic on page 1253
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261
- Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview on page 1281

 Copyright © 2013, Juniper Networks, Inc.

multiplier

Syntax multiplier *number*;

Hierarchy Level [edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the number of hello packets not received by the neighbor before Bidirectional Forwarding Detection (BFD) declares the neighbor down.

Options **number**—Maximum allowable number of hello packets missed by the neighbor.

Range: 1 through 255

Default: 3

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

nai

Syntax

```
nai (name@domain | @domain) {
  home-agent ip-address;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip dynamic-home-assignment home-agent],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip dynamic-home-assignment home-agent],
[edit routing-instances *routing-instances-name* services mobile-ip dynamic-home-assignment home-agent],
[edit services mobile-ip dynamic-home-assignment home-agent]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Configure the network address identifiers (NAI) to which registration requests are sent as part of the home agent's dynamic assignment rule .

Options

name*@*domain—User at a specified domain

@*domain*—All users at a specified domain

NOTE: The *name* can include only alphanumeric characters, dots, hyphens, or underscores. The *name* cannot end in @; @ must be used to separate *name* and *domain*. The *domain* can include only alphanumeric characters, dots, or hyphens. The *domain* must be in the format *domain.suffix*, where the *suffix* is com, org, net, and so on. The *suffix* must consist of at least two alphanumeric characters.

The remaining statement is explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring Dynamic Home Assignment for the Mobile Node on page 572

Copyright © 2013, Juniper Networks, Inc.

name-server

Syntax name-server [*server-names*];

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes]

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 168

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 47
- Configuring RADIUS Server Parameters for Subscriber Access on page 36

Copyright © 2013, Juniper Networks, Inc.

nas-port-extended-format (Access Profile)

Syntax

```
nas-port-extended-format {
   adapter-width width;
   ae-width width;
   port-width width;
   slot-width width;
   stacked;
   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.
Option **ae-width** introduced in Junos OS Release 12.1.
Option **stacked** introduced in Junos OS Release 12.3.
Option **atm** introduced in Junos OS Release 12.3R3 and supported in later 12.3R*x* releases.
Option **atm** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

Description

In an access profile, 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. You can use the same access profile to configure the NAS-Port extended format for Ethernet subscribers and ATM subscribers.

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.

vlan-width *width*—Number of bits in the VLAN ID field.

atm—Configure the NAS-Port extended format for ATM subscribers; options include:

- **adapter-width *width***—Number of bits in the adapter field.

Copyright © 2013, Juniper Networks, Inc.

- **port-width** ***width***—Number of bits in the port field.
- **slot-width** ***width***—Number of bits in the slot field.
- **vci-width** ***width***—Number of bits in the ATM virtual circuit identifier (VCI) field.
- **vpi-width** ***width***—Number of bits in the ATM virtual path identifier (VPI) field.

NOTE: Each field can be 0 through 32 bits wide; however, the total of the widths of all fields must not exceed 32 bits, or the configuration fails.

The router may truncate the values of individual fields depending on the bit width you specify.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access on page 47
- Configuring RADIUS Server Parameters for Subscriber Access on page 36

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 © 2013, 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 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65

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.

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 47
- Configuring RADIUS Server Parameters for Subscriber Access on page 36
- Configuring a NAS-Port-ID with Additional Options on page 54

nas-port-id-format (Subscriber Management)

Syntax

```
nas-port-id-format {
   agent-circuit-id;
   agent-remote-id;
   interface-description;
   nas-identifier;
}
```

Hierarchy Level

[edit access profile *profile-name* radius options]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Specify the information that the router includes in the NAS-Port-ID (RADIUS attribute 87) that it is passed to the RADIUS server during authentication and accounting. You can include any combination of the optional values.

Default

The router includes the interface description.

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.

nas-identifier—Include the NAS identifier value (RADIUS attribute 32).

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 47
- Configuring RADIUS Server Parameters for Subscriber Access on page 36
- Configuring a NAS-Port-ID with Additional Options on page 54

Copyright © 2013, Juniper Networks, Inc.

nas-port-method (Tunnel Profile)

Syntax nas-port-method cisco-avp;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Configure the LAC to interoperate with Cisco LNS devices by including the Cisco NAS Port Info AVP (100) in the ICRQ to the LNS. This AVP conveys the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

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.

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 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65

Copyright © 2013, 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
- **wireless-1x-ev**—Wireless 1xEV

Copyright © 2013, Juniper Networks, Inc.

- **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 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65

Copyright © 2013, 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.

Description Specify the port type used to authenticate subscribers. The router includes the port type in RADIUS attribute 61 (NAS-Port-Type attribute).

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

Default The router uses a port type of **ethernet**.

Options ***port-type***—One of the following port types:

- *value*—A value from 0-65535
- **adsl-cap**—Asymmetric DSL, carrierless amplitude phase (CAP) modulation
- **adsl-dmt**— Asymmetric DSL, discrete multitone (DMT)
- **async**—Asynchronous
- **cable**—Cable
- **ethernet**—Ethernet
- **fddi**—Fiber Distributed Data Interface
- **g3-fax**—G.3 Fax
- **hdlc-clear-channel**—HDLC Clear Channel
- **iapp**—Inter-Access Point Protocol (IAPP)
- **idsl**—ISDN DSL
- **isdn-sync**—ISDN Synchronous
- **isdn-v110**—ISDN Async V.110
- **isdn-v120**—ISDN Async V.120
- **piafs**—Personal Handyphone System (PHS) Internet Access Forum Standard
- **sdsl**—Symmetric DSL
- **sync**—Synchronous

Copyright © 2013, Juniper Networks, Inc.

- **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 47
- Configuring RADIUS Server Parameters for Subscriber Access on page 36

neighbor (Associate with ANCP Access Identifier)

Syntax
neighbor *ip-address*;

Hierarchy Level
[edit protocols ancp interfaces *interface-name*],
[edit protocols ancp interfaces interface-set]

Release Information
Statement introduced in Junos OS Release 9.5.

Description
Associate an ANCP neighbor with the access-loop circuit identifier (ACI) configured for the interface or interface set.

Options
ip-address—IP address of the ANCP neighbor.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring ANCP on page 1332
- Associating an Access Node with Subscribers for ANCP Operations on page 1334

Copyright © 2013, Juniper Networks, Inc.

neighbor (Define ANCP)

Syntax

```
neighbor ip-address {
    adjacency-timer seconds;
    ietf-mode;
    maximum-discovery-table-entries;
    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.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring ANCP on page 1332
- Configuring ANCP Neighbors on page 1333

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 167
- Configuring an Address-Assignment Pool for Router Advertisement on page 171

netbios-node-type

Syntax netbios-node-type *node-type*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes]

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 168

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 168

 Copyright © 2013, Juniper Networks, Inc.

network-element (Diameter Base Protocol)

Syntax

```
network-element element-name {
  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.

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 471
- Configuring Diameter Network Elements on page 473

next-hop (Dynamic Access-Internal Routes)

Syntax

next-hop *next-hop*;

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*],
[edit dynamic-profiles *profile-name* routing-options access route *prefix*]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure the next-hop address for an access route. Access routes are typically unnumbered interfaces.

Options

next-hop—Either the specific next-hop address you want to assign to the access route or one of the following next-hop address predefined variables.

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

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 690

Copyright © 2013, Juniper Networks, Inc.

next-hop-service

Syntax

```
next-hop-service {
   inside-service-interface interface-name.unit-number;
   outside-service-interface interface-name.unit-number;
   outside-service-interface-type local;
   service-interface-pool name;
}
```

Hierarchy Level

[edit services service-set *service-set-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
service-interface-pool option added in Junos OS Release 9.3.

Description

Specify interface names or a service interface pool for the forwarding next-hop service set. You cannot specify both a service interface pool and an inside or outside interface.

Options

inside-service-interface ***interface-name.unit-number***—Name and logical unit number of the service interface associated with the service set applied inside the network.

outside-service-interface ***interface-name.unit-number***—Name and logical unit number of the service interface associated with the service set applied outside the network.

outside-service-interface-type ***interface-type***—Identifies the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local.

service-interface-pool ***name***—Name of the pool of logical interfaces configured at the **[edit services service-interface-pools pool** ***pool-name*****]** hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.

NOTE: **service-interface-pool** is not applicable for IP reassembly configuration on L2TP.

Usage Guidelines

See *Configuring Service Sets to be Applied to Services Interfaces*.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

no-accounting

Syntax no-accounting;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Disable the collection of IGMP join and leave event statistics on a per-interface basis.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about disabling IGMP accounting on and interface, see "Recording IGMP Join and Leave Events" in the *Multicast Protocols Feature Guide for Routing Devices*

no-adaptation

Syntax no-adaptation;

Hierarchy Level [edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

 Copyright © 2013, Juniper Networks, Inc.

no-allow-snooped-clients

Syntax no-allow-snooped-clients;

Hierarchy Level
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name* overrides],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay group *group-name* interface *interface-name* overrides],
[edit forwarding-options dhcp-relay group *group-name* overrides],
[edit forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay ...],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay ...]

Release Information
Statement introduced in Junos OS Release 10.2.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description
Explicitly disable DHCP snooping support on the router or switch.

Use the statement at the **[edit ... dhcpv6]** hierarchy levels to explicitly disable snooping support on the router or switch for DHCPv6 relay agent.

NOTE: In Junos OS Release 10.0 and earlier, DHCP snooping is *enabled* by default. In Release 10.1 and later, DHCP snooping is *disabled* by default.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Extended DHCP Relay Agent Overview on page 276
- Overriding the Default DHCP Relay Configuration Settings on page 293
- DHCP Snooping Support on page 299

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 © 2013, Juniper Networks, Inc.

Related Documentation

- Extended DHCP Relay Agent Overview on page 276
- Overriding the Default DHCP Relay Configuration Settings on page 293
- Disabling Automatic Binding of Stray DHCP Requests on page 320

no-keepalives (Dynamic Profiles)

Syntax

no-keepalives;

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit"]

Release Information

Statement introduced before Junos OS Release 7.4.
Support of the **[edit dynamic-profiles *profile-name*]** hierarchy level introduced in Junos OS Release 9.5.
Support of the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit"]** hierarchy level introduced in Junos OS Release 10.1.

Description

Disable the sending of keepalives.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring Dynamic Authentication for PPP Subscribers on page 365

no-qos-adjust (Dynamic Routing Options)

Syntax

no-qos-adjust;

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-options multicast interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Disable hierarchical bandwidth adjustment for all dynamically created subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Multicast with Subscriber VLANs*

no-vlan-id-validate

Syntax no-vlan-id-validate;

Hierarchy Level [edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name*],
[edit protocols l2circuit neighbor *address* interface *interface-name*]

Release Information Statement introduced in Junos OS Release 13.1.

Description Uniquely identify a Layer 2 circuit for either a standard pseudowire or a redundant pseudowire.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Interfaces for Layer 2 Circuits*
- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Configuring Layer 2 Circuit Signaling for Pseudowire Subscriber Logical Interfaces on page 937

oam-on-svlan (Ethernet Interfaces)

Syntax oam-on-svlan;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 13.1.

Description On MX Series routers with MPC/MIC interfaces, enable propagation of the Ethernet IEEE 802.1ag Operation, Administration, and Maintenance (OAM) state of a static single-tagged service VLAN (S-VLAN) logical interface to the dynamic or static double-tagged customer VLAN (C-VLAN) logical interface and associated subscriber interfaces configured on the S-VLAN. The static S-VLAN logical interface must be configured with Ethernet OAM connectivity fault management (CFM) on a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet physical interface. The C-VLAN logical interface must have the same S-VLAN (outer) tag as the S-VLAN logical interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Support for Service VLANs with Double-Tagged Customer VLANs on page 743
- Ethernet OAM Support for Service VLANs Overview on page 709

 Copyright © 2013, Juniper Networks, Inc.

oif-map (Dynamic IGMP Interface)

Syntax oif-map *map-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Associates an OIF map to the IGMP interface using a dynamic profile. The OIF map is a routing policy statement that can contain multiple terms.

Options ***map-name***—Name of the OIF map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Multicast with Subscriber VLANs*
- Configuring a Dynamic Profile for Client Access on page 679

oif-map (Dynamic MLD Interface)

Syntax oif-map *map-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Associate an outgoing interface (OIF) map to a dynamic MLD logical interface. The OIF map is a routing policy statement that can contain multiple terms.

Options ***map-name***—Name of the OIF map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Multicast with Subscriber VLANs*

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 Allocates and de-allocates an IPv4 address after initial PPP authentication for a subscriber who does not have an existing IPv4 address and can be configured at either the interface level or at the system level. Disabled by default. When configured at the interface level, dynamic profile changes take effect only for any new subscriber logins. Changes for static PPP IFLs logs out the subscriber. When configured at the system level, globally enables an on-demand-ip-address for PPP subscribers. If configured at both the interface level and the system level, the system level configuration takes precedence and changes take effect only for new subscriber logins.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Saving IPv4 Addresses for Dual-Stack PPP Subscribers*

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 © 2013, 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]

Release Information Statement introduced in Junos OS Release 9.0.

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, 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 168

option-18 (Interface-ID for DHCPv6 Autosense VLANs)

Syntax option-18;

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges authentication username-include], [edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication username-include]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify that Option 18 (Interface-ID) information received in the innermost DHCPv6 Relay-Forward message header 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

- Configuring VLAN Interface Username Information for AAA Authentication on page 729
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 329
- Creating Unique Usernames for DHCP Clients on page 239
- Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs on page 731
- **option-37 (Relay Agent Remote-ID for DHCPv6 Autosense VLANs) on page 1825**

 Copyright © 2013, Juniper Networks, Inc.

option-37 (Relay Agent Remote-ID for DHCPv6 Autosense VLANs)

Syntax option-37;

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges authentication username-include], [edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication username-include]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify that Option 37 (DHCPv6 Relay Agent Remote-ID) information, received in the innermost DHCPv6 Relay-Forward message header, 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

- Configuring VLAN Interface Username Information for AAA Authentication on page 729
- **relay-agent-remote-id on page 1932**
- Creating Unique Usernames for DHCP Clients on page 239
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 329
- Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs on page 731
- **option-18 (Interface-ID for DHCPv6 Autosense VLANs) on page 1824**

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 214

Copyright © 2013, 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 214

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]

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.

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 168

Copyright © 2013, 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 214

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 215
- Extended DHCP Local Server Overview on page 200
- Address-Assignment Pools Overview on page 167

Copyright © 2013, 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 214

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]

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.

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 168

Copyright © 2013, 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 323
- *Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)*

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;
  }
  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:
      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;
    nas-identifier;
  }
  nas-port-type {
    ethernet {
      port-type;
    }
  }
  revert-interval interval;
  vlan-nas-port-stacked-format;
}
```


Copyright © 2013, Juniper Networks, Inc.

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.

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 47
- RADIUS Server Options for Subscriber Access on page 41

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 24

order (Mobile IP)

Syntax

order (aaa | local);

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip authenticate],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip authenticate],
[edit routing-instances *routing-instances-name* services mobile-ip authenticate],
[edit services mobile-ip authenticate]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the authentication method performed for Mobile IP.

Default

AAA is the default authentication method.

Options

aaa—Authentication is performed by AAA. This option is available only in the default router and default routing instance, and therefore only in the **[edit services mobile-ip]** hierarchy level.

local—Authentication is performed using parameters defined in the local database.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Access Type for Mobile IP on page 573

Copyright © 2013, Juniper Networks, Inc.

origin (Diameter Base Protocol)

Syntax

```
origin {
  host hostname;
  realm realm-name;
}
```

Hierarchy Level [edit diameter]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify values of Origin-Realm-AVP and Origin-Host-AVP used in all messages sent by the Diameter instance.

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

- Configuring Diameter on page 471
- Configuring the Origin Attributes of the Diameter Instance on page 472

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

- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

output (Dynamic Service Sets)

Syntax

```
service-set service-set-name {
    service-filter filter-name;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* service],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos–interface–unit" family *family* service]

Release Information

Statement introduced in Junos OS Release 9.5.
Support of the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* service]** hierarchy level introduced in Junos OS Release 10.1.

Description

Define the output service sets and filters to be applied to traffic by a dynamic profile. Only the Internet Protocol version 4 (IPv4) protocol family is currently supported for dynamic PPPoE logical interfaces.

The remaining statement is explained separately.

Options

service-set-name—Name of the service set.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Service Sets Overview on page 1151
- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197

Copyright © 2013, Juniper Networks, Inc.

output-traffic-control-profile (Dynamic CoS Definition)

Syntax output-traffic-control-profile (*profile-name* | $junos-cos-traffic-control-profile);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.
Variable **$junos-cos-traffic-control-profile** introduced in Junos OS Release 11.2.

Description Apply an output traffic scheduling and shaping profile to the logical interface.

Options *profile-name*—Name of the traffic-control profile to be applied to this interface

$junos-cos-traffic-control-profile—Variable for the traffic-control profile that is specified for the logical interface. The variable is replaced with the traffic-control profile when the subscriber is authenticated at login.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile on page 979
- Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers on page 1125
- **traffic-control-profiles on page 2055**

output-vlan-map (Dynamic Interfaces)

Syntax

```
output-vlan-map {
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  (pop | swap);
  tag-protocol-id tpid;
  vlan-id number;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

For dynamic interfaces, define the rewrite profile to be applied to outgoing frames on this logical interface.

The 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

- *Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution*

Copyright © 2013, Juniper Networks, Inc.

overhead-accounting (Dynamic Traffic Shaping)

Syntax

```
overhead-accounting {
  bytes bytes;
  cell-mode cell-mode-bytes cell-mode-bytes;
  frame-mode frame-mode-bytes frame-mode-bytes;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Configure the mode to shape downstream ATM traffic based on either frames or cells.

Default

The default is frame-mode.

Options

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

- CoS Adjustment Control Profiles Overview on page 1084
- Configuring CoS Adjustment Control Profiles on page 1106
- **adjustment-control-profiles on page 1442**
- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 1092
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072
- *egress-shaping-overhead*
- **bytes on page 1488**
- **cell-mode on page 1494**
- **frame-mode on page 1651**

overrides (DHCP Local Server)

Syntax

```
overrides {
  client-discover-match;
  delegated-pool;
  interface-client-limit number;
  process-inform {
    pool pool-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.

Description

Override the default configuration settings for the extended DHCP local server. Specifying the **overrides** statement with no subordinate statements removes all DHCP local server overrides at that hierarchy level.

- To override global DHCP local server configuration options, include the **overrides** statement and its subordinate statements at the **[edit system services dhcp-local-server]** hierarchy level.
- To override configuration options for a named group of interfaces, include the statements at the **[edit system services dhcp-local-server group *group-name*]** hierarchy level.
- To override configuration options for a specific interface within a named group of interfaces, include the statements at the **[edit system services dhcp-local-server group *group-name* interface *interface-name*]** hierarchy level.
- Use the **[edit system services dhcp-local-server dhcpv6]** hierarchy level to override DHCPv6 configuration options.

The remaining statements are explained separately.

The interface-client-limit statement id not supported in the **[edit system services dhcp-local-server dhcpv6]** hierarchy level.

The delegated-pool and the rapid-commit statements are supported in the **[edit system services dhcp-local-server dhcpv6 ...]** hierarchy level only.

Copyright © 2013, 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

- Extended DHCP Local Server Overview on page 200
- Overriding Default DHCP Local Server Configuration Settings on page 220
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

overrides (DHCP Relay Agent)

Syntax

```
overrides {
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  client-discover-match <option60–and-option82>;
  disable-relay;
  interface-client-limit number;
  layer2-unicast-replies;
  no-allow-snooped-clients;
  no-bind-on-request;
  proxy-mode;
  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 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 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. All other statements are supported at the **dhcp-relay** hierarchy levels only.

- **allow-snooped-clients**
- **interface-client-limit**
- **no-allow-snooped-clients**
- **no-bind-on-request**
- **send-release-on-delete**

Copyright © 2013, Juniper Networks, Inc.

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 276
- Overriding the Default DHCP Relay Configuration Settings on page 293
- Deleting DHCP Local Server and DHCP Relay Override Settings on page 229
- dhcp-relay on page 1548

packet-triggered-subscribers

Syntax

```
packet-triggered-subscribers {
  partition partition-name {
    destination-host hostname;
    destination-realm realm;
    diameter-instance instance-name;
  }
  traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
  }
}
```

Hierarchy Level

[edit system services]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Configure PTSP to interact with an SAE in an SRC environment to provision packet-triggered subscribers.

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

- Configuring the PTSP Partition on page 526

packet-triggered-subscribers-partition

Syntax packet-triggered-subscribers-partition *partition-name*;

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the PTSP partition to associate with the logical system and routing instance.

Options ***partition-name***—Name of the PTSP partition that you want PTSP to use. The name is defined with the **partition** statement at the **[edit system services packet-triggered-subscribers]** hierarchy level.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Assigning the PTSP Partition on page 526

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.

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 190

 Copyright © 2013, Juniper Networks, Inc.

pap

Syntax

```
pap {
   access-profile name;
   default-pap-password password;
   local-name name;
   local-password password;
   passive;
}
```

Hierarchy Level

[edit interfaces *interface-name* ppp-options],
[edit interfaces *interface-name* unit *logical-unit-number* ppp-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options]

Release Information

Statement introduced in Junos OS Release 8.3.

Description

Configure the Password Authentication Protocol (PAP). Use PAP authentication as a means to provide a simple method for the peer to establish its identity using a two-way handshake. This is done only upon initial link establishment.

After the link is established, an ID and password pair is repeatedly sent by the peer to the authenticator until authentication is acknowledged or the connection is terminated.

BEST PRACTICE: **On inline service (si) interfaces for L2TP, only the pap statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.**

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Configuring PPP PAP Authentication*
- *Tracing Operations of the pppd Process*
- *traceoptions (PPP Process)*
- *Junos OS Administration Library for Routing Devices*
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 419

Copyright © 2013, Juniper Networks, Inc.

pap (Dynamic PPP)

Syntax pap;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options], [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]** hierarchy level introduced in Junos OS Release 12.2.

Description Specify PAP authentication in a PPP dynamic profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring Dynamic Authentication for PPP Subscribers on page 365
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 370
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 419

pap (L2TP)

Syntax pap;

Hierarchy Level [edit access group-profile *profile-name* ppp ppp-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description (MX Series routers only) Specify PAP authentication for PPP subscribers in an L2TP LNS user group profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 418

 Copyright © 2013, 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 right-to-left

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 Names on page 188
- Configuring Domain Name Usage for Domain Maps on page 187

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.

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 491
- Configuring the JSRC Partition on page 492

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.

Description Configure a Gx-Plus partition.

Options *partition-name*—Name of the Gx-Plus partition.

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

- Configuring Gx-Plus on page 549
- Configuring the Gx-Plus Partition on page 550

Copyright © 2013, Juniper Networks, Inc.

partition (PTSP)

Syntax

```
partition partition-name {
   destination-host hostname;
   destination-realm realm;
   diameter-instance instance-name;
}
```

Hierarchy Level [edit system services packet-triggered-subscribers]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure a PTSP partition.

Options *partition-name*—Name of the PTSP partition.

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

- Configuring the PTSP Partition on page 526

passive (Dynamic IGMP Interface)

Syntax passive <allow-receive> <send-general-query> <send-group-query>;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6.
allow-receive, **send-general-query**, and **send-group-query** options were introduced in Junos OS Release 10.0.

Description Dynamically specify that IGMP run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as IGMP reports, queries, and leaves.

NOTE: You can selectively activate up to two out of the three available options for the passive statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the passive statement.

Options **allow-receive**—(Optional) Enables IGMP to receive control traffic on the interface.

send-general-query—(Optional) Enables IGMP to send general queries on the interface.

send-group-query—(Optional) Enables IGMP to send group-specific and group-source-specific queries on the interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Multicast with Subscriber VLANs*
- For general information about configuring IGMP, see the *Multicast Protocols Feature Guide for Routing Devices*.

 Copyright © 2013, Juniper Networks, Inc.

passive (Dynamic MLD Interface)

Syntax passive <allow-receive> <send-general-query> <send-group-query>;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Specify that MLD run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as MLD reports, queries, and leaves.

NOTE: You can selectively activate up to two out of the three available options for the passive statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the passive statement.

Options **allow-receive**—(Optional) Enables MLD to receive control traffic on the interface.

send-general-query—(Optional) Enables MLD to send general queries on the interface.

send-group-query—(Optional) Enables MLD to send group-specific and group-source-specific queries on the interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Multicast with Subscriber VLANs*

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.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214

Copyright © 2013, 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 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.

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.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214
- Configuring Passwords for Usernames on page 238

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.

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 500
- Configuring the Static Subscriber Global Authentication Password on page 503
- Configuring the Static Subscriber Group Authentication Password on page 506

Copyright © 2013, Juniper Networks, Inc.

peak-burst-size

Syntax peak-burst-size *bytes*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall three-color-policer *name* two-rate],
[edit firewall three-color-policer *policer-name* two-rate]

Release Information Statement introduced in Junos OS Release 7.4.
Support at the **[edit dynamic-profiles ... two-rate]** hierarchy level introduced in Junos OS Release 11.4.

Description For a two-rate three-color policer, configure the peak burst size (PBS) as a number of bytes. The PBS defines the maximum number of bytes of unused peak bandwidth capacity that can be accumulated. The accumulated bandwidth allows for moderate periods of bursting traffic that exceeds the peak information rate (PIR) and the committed burst size (CBS).

NOTE: When you include the **peak-burst-size** statement in the configuration, you must also include the **committed-burst-size** and **peak-information-rate** statements at the same hierarchy level.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits.

- A traffic flow is categorized green if it conforms to both the committed information rate (CIR) and the CBS-bounded accumulation of available committed bandwidth capacity.
- A traffic flow is categorized yellow if exceeds the CIR and CBS but conforms to the PIR. Packets in a yellow flow are marked with **medium-high** packet loss priority (PLP) and then passed through the interface.
- A traffic flow is categorized red if exceeds the PIR and the PBS-bounded accumulation of available peak bandwidth capacity. Packets in a red traffic flow are marked with **high** PLP and then either passed through the interface or optionally discarded.

Options *bytes*—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 100,000,000,000 bytes

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation
- *Three-Color Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*

- *Policer Color-Marking and Actions*
- *Dual Token Bucket Algorithms*
- *Determining Proper Burst Size for Traffic Policers*
- committed-burst-size on page 1516
- committed-information-rate on page 1518
- excess-burst-size on page 1619
- peak-information-rate on page 1859

Copyright © 2013, Juniper Networks, Inc.

peak-information-rate

Syntax peak-information-rate *bps*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall three-color-policer *name* two-rate],
[edit firewall three-color-policer *policer-name* two-rate]

Release Information Statement introduced in Junos OS Release 7.4.
Support at the **[edit dynamic-profiles ... two-rate]** hierarchy level introduced in Junos OS Release 11.4.

Description For a two-rate three-color policer, configure the peak information rate (PIR) as a number of bits per second. The PIR is the maximum rate for traffic arriving at or departing from the interface under peak line conditions. Traffic that exceeds the committed information rate (CIR) and the committed burst size (CBS) is metered to the PIR.

NOTE: When you include the **peak-information-rate** statement in the configuration, you must also include the **committed-information-rate** and **peak-burst-size** statements at the same hierarchy level.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits.

- A traffic flow is categorized green if it conforms to both the CIR and the CBS-bounded accumulation of available committed bandwidth capacity.
- A traffic flow is categorized yellow if exceeds the CIR and CBS but conforms to the PIR. Packets in a yellow flow are marked with **medium-high** packet loss priority (PLP) and then passed through the interface.
- A traffic flow is categorized red if exceeds the PIR and the PBS-bounded accumulation of available peak bandwidth capacity. Packets in a red traffic flow are marked with **high** PLP and then either passed through the interface or optionally discarded.

Options *bps*—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 100,000,000,000 bps

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Three-Color Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Dual Token Bucket Algorithms*

Copyright © 2013, Juniper Networks, Inc.

- *Determining Proper Burst Size for Traffic Policers*
- committed-burst-size on page 1516
- committed-information-rate on page 1518
- excess-burst-size on page 1619
- peak-burst-size on page 1857

Copyright © 2013, Juniper Networks, Inc.

peer

Syntax

```
peer {
  (ip-address address | nai name@domain) {
    spi hexadecimal-value {
      algorithm (hmac-md5 | md5);
      entity-type (host | mobility-agent);
      key (hex | ascii) string;
      replay-method (timestamp seconds | none);
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip],
[edit routing-instances *routing-instances-name* services mobile-ip],
[edit services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the authentication configurations for a home agent mobile node. An authentication enables the registration message as acceptable to the final recipient of the registration message.

Options

ip-address *address*—IP address of the peer.

nai *name*@*domain*—Network address identifier (NAI) of the peer. The *name* can include only alphanumeric characters, dots, hyphens, or underscores. The *name* cannot end in @; @ must be used to separate *name* and *domain*. The *domain* can include only alphanumeric characters, dots, or hyphens The *domain* must be in the format *domain.suffix*, where the *suffix* is com, org, net, and so on. The *suffix* must consist of at least two alphanumeric characters.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

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>;
  routing-instance routing-instance-name;
}
```

Hierarchy Level [edit diameter]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a remote peer for the Diameter instance.

Options *peer-name*—Name of the peer.

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

- Configuring Diameter on page 471
- Configuring Diameter Peers on page 473

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 Define and prioritize a peer associated with a Diameter network element.

Options *peer-name*—Name of the peer.

The remaining statement is explained separately.

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 471
- Configuring Diameter Network Elements on page 473

 Copyright © 2013, Juniper Networks, Inc.

physical-interface-policer

Syntax

physical-interface-policer;

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall policer *policer-name*],
[edit firewall policer *policer-name*],
[edit firewall three-color-policer *policer-name*],
[edit logical-system *logical-system-name* firewall policer *policer-name*],
[edit logical-system *logical-system-name* three-color-policer *policer-name*],
[edit routing-instances *routing-instance-name* firewall policer *policer-name*],
[edit routing-instances *routing-instance-name* firewall three-color-policer *policer-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* firewall policer *policer-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* firewall three-color-policer *policer-name*]

Release Information

Statement introduced in Junos OS Release 9.6.
Support at the **[edit dynamic-profiles ... policer *policer-name*]** hierarchy level introduced in Junos Release OS 11.4.

Description

Configure an aggregate policer for a physical interface.

A physical interface policer can be a two-color or three-color policer. When you apply physical interface policer, to different protocol families on the same logical interface, the protocol families share the same policer instance. This means that rate limiting is performed aggregately for the protocol families for which the policer is applied. This feature enables you to use a single policer instance to perform aggregate policing for different protocol families on the same physical interface. If you want a policer instance to be associated with a protocol family, the corresponding physical interface filter needs to be applied to that protocol family. The policer is not automatically applied to all protocol families configured on the physical interface.

In contrast, with logical interface policers there are multiple separate policer instances.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Two-Color and Three-Color Physical Interface Policers*
- *physical-interface-filter*

pic (M Series, MX Series, and T Series Routers)

Syntax

```
pic pic-number {
  ce1 {
    e1 port-number {
      channel-group group-number timeslots slot-number;
    }
  }
  ct3 {
    port port-number {
      t1 link-number {
        channel-group group-number timeslots slot-number;
      }
    }
  }
  framing (sdh | sonet);
  idle-cell format {
    itu-t;
    payload-pattern payload-pattern-byte;
  }
  inline-services {
    bandwidth (1g | 10g);
  }
  max-queues-per-interface (8 | 4);
  no-concatenate;
}
```

Hierarchy Level

[edit chassis fpc *slot-number*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure properties for an individual PIC.

Options

pic-number—Slot number in which the PIC is installed.

Range: 0 through 3

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Junos OS to Enable SONET/SDH Framing for SONET/SDH PICs*
- *Configuring the Junos OS to Enable a SONET PIC to Operate in Channelized (Multiplexed) Mode*
- *Configuring the Junos OS to Support Channelized DS3-to-DS0 Naming for Channel Groups and Time Slots*
- *Configuring the Junos OS to Support Channel Groups and Time Slots for Channelized E1 PICs*
- Enabling Inline Service Interfaces on page 427

Copyright © 2013, Juniper Networks, Inc.

policer (Configuring)

Syntax

```
policer policer-name {
   filter-specific;
   if-exceeding {
      bandwidth-limit bps;
      bandwidth-percent number;
      burst-size-limit bytes;
   }
   logical-bandwidth-policer;
   logical-interface-policer;
   physical-interface-policer;
   shared-bandwidth-policer;
   then {
      policer-action;
   }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall],
[edit firewall],
[edit logical-systems *logical-system-name* firewall]

Release Information

Statement introduced before Junos OS Release 7.4.
The **out-of-profile** policer action added in Junos OS Release 8.1.
The **logical-bandwidth-policer** statement added in Junos OS Release 8.2.
Logical systems support introduced in Junos OS Release 9.3.
The **physical-interface-policer** statement introduced in Junos OS Release 9.6.
The **shared-bandwidth-policer** statement added in Junos OS Release 11.2
Support at the **[edit dynamic-profiles ... firewall]** hierarchy level introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure policer rate limits and actions. When included at the **[edit firewall]** hierarchy level, the **policer** statement creates a template, and you do not have to configure a policer individually for every firewall filter or interface. To activate a policer, you must include the **policer-action** modifier in the **then** statement in a firewall filter term or on an interface.

Options

policer-action—One or more actions to take:

- **discard**—Discard traffic that exceeds the rate limits.
- **forwarding-class** *class-name*—Specify the particular forwarding class.
- **loss-priority**—Set the packet loss priority (PLP) to **low**, **medium-low**, **medium-high**, or **high**.
- **out-of-profile**—On J Series routers with strict priority queuing, prevent starvation of other queues by rate limiting the data stream entering the strict priority queue, marking the packets that exceed the rate limit as out-of-profile, and dropping the out-of-profile packets if the physical interface is congested.

Copyright © 2013, Juniper Networks, Inc.

policer-name—Name that identifies the policer. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”). Policer names cannot begin with an underscore in the form **__.***.

then—Actions to take on matching packets.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Bandwidth Policer Overview*
- *Configuring Firewall Filters and Policers for VPLS*
- *Configuring Multifield Classifiers*
- *Logical Interface (Aggregate) Policer Overview*
- *Physical Interface Policer Overview*
- *Statement Hierarchy for Configuring Policers*
- *Single-Rate Two-Color Policer Overview*
- *Using Multifield Classifiers to Set PLP*
- filter (Configuring) on page 1639
- *priority (Schedulers)*

Copyright © 2013, Juniper Networks, Inc.

policy (Subscriber Secure Policy)

Syntax

```
policy policy-name {
  inet {
    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;
      }
    }
  }
  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;
      }
    }
  }
}
```

Hierarchy Level [edit services radius-flow-tap]

Release Information Statement introduced in Junos OS Release 12.3.

Description Specify the policy that is applied to mirrored packets sent to a mediation device.

Options ***policy-name***—Name of the policy from which to drop traffic.

The remaining statements are explained separately.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

policy-options

Syntax

```
policy-options {
  as-path name regular-expression;
  as-path-group group-name;
  community name {
    invert-match;
    members [ community-ids ];
  }
  condition condition-name {
    if-route-exists address table table-name;
  }
  damping name {
    disable;
    half-life minutes;
    max-suppress minutes;
    reuse number;
    suppress number;
  }
  policy-statement policy-name {
    term term-name {
      from {
        family;
        match-conditions;
        policy subroutine-policy-name;
        prefix-list name;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
      }
      to {
        match-conditions;
        policy subroutine-policy-name;
      }
      then actions;
      default-action (accept | reject);
    }
  }
  prefix-list name {
    ip-addresses;
  }
}
```

Hierarchy Level

[edit],
[edit dynamic],
[edit dynamic-profiles *profile-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Support at the **[edit dynamic-profiles]** hierarchy level introduced in Junos OS Release 11.4.

Description

Configure routing policy.

The statements are explained separately.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Using Routing Policy in an ISP Network*

pool (Address-Assignment Pools)

Syntax

```
pool pool-name {
  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;
    }
  }
  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.

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 167
- Configuring Address-Assignment Pools on page 168
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

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 © 2013, 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 227
- Overriding Default DHCP Local Server Configuration Settings on page 220

pool (L2TP Service Interfaces)

Syntax

```
pool pool-name {
  interface service-interface-name;
}
```

Hierarchy Level [edit services service-device-pools]

Release Information Statement introduced in Junos OS Release 11.4.

Description Define a pool of service interfaces that can be assigned to an L2TP tunnel group for traffic load-balancing. The service device pool is required for dynamic LNS sessions.

Options *pool-name*—Name of the service interface pool.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 430

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.

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 215
- Extended DHCP Local Server Overview on page 200
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

Copyright © 2013, Juniper Networks, Inc.

pop (Dynamic VLANs)

Syntax pop;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 10.4.

Description For dynamic VLAN interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Removing a VLAN Tag*
- *Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution*

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 23
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24

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 471
- Configuring Diameter Peers on page 473

post-service-filter (Dynamic Service Sets)

Syntax post-service-filter *filter-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* service input],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos–interface–unit" family *family* service input]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* service input]** hierarchy level introduced in Junos OS Release 10.1.

Description Define the filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected. You can configure a postservice filter on the input side of the interface only. Only the Internet Protocol version 4 (IPv4) protocol family is currently supported for dynamic PPPoE logical interfaces.

Options ***filter-name***—Identifier for the post-service filter.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Dynamic Service Sets Overview on page 1151
- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197

 Copyright © 2013, Juniper Networks, Inc.

pp0 (Dynamic PPPoE)

Syntax

```
pp0 {
  unit logical-unit-number {
    keepalives interval seconds;
    no-keepalives;
    pppoe-options {
      underlying-interface interface-name;
      server;
    }
    ppp-options {
      authentication [ authentication-protocols ];
      chap {
        challenge-length minimum minimum-length maximum maximum-length;
      }
      pap;
    }
    family inet {
      unnumbered-address interface-name;
      address address;
      service {
        input {
          service-set service-set-name {
            service-filter filter-name;
          }
          post-service-filter filter-name;
        }
        output {
          service-set service-set-name {
            service-filter filter-name;
          }
        }
      }
      filter {
        input filter-name {
          precedence precedence;
        }
        output filter-name {
          precedence precedence;
        }
      }
    }
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the dynamic PPPoE logical interface in a dynamic profile. When the router creates a dynamic PPPoE logical interface on an underlying Ethernet interface configured with PPPoE (**ppp-over-ether**) encapsulation, it uses the information in the dynamic profile to determine the properties of the dynamic PPPoE logical interface.

Copyright © 2013, Juniper Networks, Inc.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Basic PPPoE Dynamic Profile on page 900
- Configuring a PPPoE Dynamic Profile with Additional Options on page 903
- Configuring Dynamic Authentication for PPP Subscribers on page 365
- For information about creating static PPPoE interfaces, see *Configuring PPPoE*

ppp (Group Profile)

Syntax

```
ppp {
   cell-overhead;
   encapsulation-overhead bytes;
   framed-pool framed-pool;
   idle-timeout seconds;
   interface-id interface-id;
   keepalive seconds;
   ppp-options {
      chap;
      pap;
   }
   primary-dns primary-dns;
   primary-wins primary-wins;
   secondary-dns secondary-dns;
   secondary-wins secondary-wins;
}
```

Hierarchy Level

[edit access group-profile *profile-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure PPP properties for a group profile.

The remaining statements are explained separately.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 418

Copyright © 2013, Juniper Networks, Inc.

pppoe-options (Dynamic PPPoE)

Syntax

```
pppoe-options {
   underlying-interface interface-name;
   server;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit"]

Release Information

Statement introduced in Junos OS Release 10.1.

Description

Configure the underlying interface and PPPoE server mode for a dynamic PPPoE logical interface in a dynamic profile.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Basic PPPoE Dynamic Profile on page 900
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

pppoe-tags (Adjustment Control Profiles)

Syntax

```
pppoe-tags {
   priority priority;
   algorithm algorithm;
}
```

Hierarchy Level

[edit class-of-service adjustment-control-profiles *profile-name* application]

Release Information

Statement introduced in Junos OS Release 13.1.

Description

Configure the shaping rate adjustment controls for the Point-to-Point Protocol over Ethernet (PPPoE) Tags application.

Options

priority—Priority of the Point to Point Protocol over Ethernet IA Tags application in the adjustment control profile.

Range: 1 through 10; 1 being the highest priority.

Default: 2

algorithm—Rate adjustment algorithm used by the Point to Point Protocol over Ethernet (PPPoE) IA Tags application.

Values:

- adjust-never—Do not perform rate adjustments.
- adjust-always—Adjust the shaping rate unconditionally.
- adjust-less—Adjust the shaping rate if it is less than the configured value.
- adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
- adjust-greater—Adjust the shaping rate if it is greater than the configured value.
- adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.

Default: adjust-less

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- CoS Adjustment Control Profiles Overview on page 1084
- Configuring CoS Adjustment Control Profiles on page 1106
- Verifying the CoS Adjustment Control Profile Configuration on page 1107
- adjustment-control-profiles on page 1442
- application (Adjustment Control Profiles) on page 1459

Copyright © 2013, Juniper Networks, Inc.

pppoe-underlying-options (Dynamic VLAN Interface Sets)

Syntax

```
pppoe-underlying-options {
  max-sessions number;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces interface-set "$junos-interface-set-name"]

Release Information

Statement introduced in Junos OS Release 12.2.

Description

Configure PPPoE-specific interface properties in the dynamic profile that defines the agent circuit identifier (ACI) interface set. An ACI interface set is a logical collection of subscriber interfaces that originate at the same household or on the same access-loop port. Configuring PPPoE-specific interface properties for an ACI interface set enables you to apply these attributes to all subscribers on a per-household basis.

The remaining statement is explained separately.

NOTE: When you configure PPPoE-specific interface properties for an ACI interface set, only the max-sessions statement is currently supported.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic VLANs Based on Agent Circuit Identifier Information on page 733
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704

Copyright © 2013, Juniper Networks, Inc.

pppoe-underlying-options (Static and Dynamic Subscribers)

Syntax

```
pppoe-underlying-options {
  access-concentrator name;
  dynamic-profile profile-name;
  duplicate-protection;
  max-sessions number;
  max-sessions-vsa-ignore;
  service-name-table table-name;
  short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Configure PPPoE-specific interface properties for the underlying interface on which the router creates a static or dynamic PPPoE logical interface. The underlying interface must be configured with PPPoE (**ppp-over-ether**) encapsulation.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring PPPoE* (for static interfaces)
- Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces on page 905
- *Assigning a Service Name Table to a PPPoE Underlying Interface*

Copyright © 2013, Juniper Networks, Inc.

ppp-options

Syntax

```
ppp-options {
  authentication [ authentication-protocols ];
  chap {
    access-profile name;
    challenge-length minimum minimum-length maximum maximum-length;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
  dynamic-profile profile-name;
  lcp-max-conf-req number
  lcp-restart-timer milliseconds;
  loopback-clear-timer seconds;
  ncp-max-conf-req number
  ncp-restart-timer milliseconds;
  on-demand-ip-address
  pap {
    access-profile name;
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

On interfaces with PPP encapsulation, configure PPP-specific interface properties.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

BEST PRACTICE: On inline service (si) interfaces for L2TP, only the **chap** and **pap** statements are typically used for subscriber management. We recommend that you leave the other statements subordinate to **ppp-options**—including those subordinate to **chap** and **pap**—at their default values.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 419

ppp-options (Dynamic PPP)

Syntax

```
ppp-options {
  authentication [ authentication-protocols ];
  chap {
    challenge-length minimum minimum-length maximum maximum-length;
  }
  on-demand-ip-address;
  pap;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit”],
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]** hierarchy level introduced in Junos OS Release 12.2.

Description

Configure PPP-specific interface properties in a dynamic profile.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Profiles Overview on page 638
- Configuring Dynamic Authentication for PPP Subscribers on page 365
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 370
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 419

Copyright © 2013, Juniper Networks, Inc.

ppp-options (L2TP)

Syntax

```
ppp-options {
  chap;
  pap;
}
```

Hierarchy Level

[edit access group-profile *profile-name* ppp]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Configure PPP-specific properties in a group profile that applies to tunneled PPP subscribers at the LNS.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 418

ppp-subscriber-services

Syntax ppp-subscriber-services (disable | enable);

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 10.2.

Description Enable dynamic PPP subscriber services on non-PPPoE interfaces on certain PICs.

NOTE: When you include this statement, the relevant PICs restart. This action disrupts subscribers already logged in through those PICs. You can confirm completion of the restart by issuing the **show chassis pic fpc-slot *slot-number* pic-slot *slot-number*** command.

Options **disable**—Disable subscriber services.

enable—Enable subscriber services.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *show chassis pic*
- Attaching Dynamic Profiles to MLPPP Bundles on page 375
- For hardware requirements, see Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces on page 374

 Copyright © 2013, Juniper Networks, Inc.

precedence

Syntax precedence *precedence*;

Hierarchy Level
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter input *filter-name*],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter output *filter-name*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family* filter input *filter-name*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family* filter output *filter-name*],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos–interface–unit" family *family* filter input *filter-name*],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos–interface–unit" family *family* filter output *filter-name*]

Release Information Statement introduced in Junos OS Release 9.3.
The **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family inet filter input *filter-name*]** hierarchy level and **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family inet filter output *filter-name*]** hierarchy level introduced in Junos OS Release 10.1.

Description Apply a precedence to a dynamic filter. Only the Internet Protocol version 4 (IPv4) protocol family is currently supported for dynamic PPPoE logical interfaces.

Options ***precedence***—Precedence value for the filter. The lower the precedence value, the higher the precedence.

Range: 0 through 250

Default: 0

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- For general information about configuring firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*
- Dynamic Firewall Filters Overview on page 1132
- Classic Filters Overview on page 1133
- Fast Update Filters Overview on page 1146
- Basic Classic Filter Syntax on page 1135
- Basic Fast Update Filter Syntax on page 1149

Copyright © 2013, Juniper Networks, Inc.

predefined-variable-defaults (Dynamic Profiles)

Syntax predefined-variable-defaults *predefined-variable* <*variable-option*> *default-value*

Hierarchy Level [edit dynamic-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure default values for the predefined variables that are configured in a dynamic profile. These default values are used when RADIUS does not supply a value.

Options

predefined-variable—Name of the predefined variable to which you want to assign a default value. Do not include the junos prefix.

variable-option—Name of the specific variable option to which you want to assign a default value.

default-value—Default value that you want to assign to the predefined variable. Only certain predefined variables support multiple default values.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- For a list of predefined variables and options for which you can configure default values, see Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs on page 660
- Configuring Default Values for Predefined Variables in a Dynamic Profile on page 675

 Copyright © 2013, Juniper Networks, Inc.

preference (Subscriber Management)

Syntax preference *route-distance*

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*],
[edit dynamic-profiles *profile-name* routing-options access route *prefix*]

Release Information Statement introduced in Junos OS Release 9.5.
Support at **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*]** hierarchy levels introduced in Junos OS Release 10.1.

Description Dynamically configure the distance for an access route.

Options ***route-distance***—Either the specific distance you want to assign to the access route or either of the following distance variables:

- **$junos-framed-route-distance**—Distance of an IPv4 access route; the variable is dynamically replaced with the preference value (Subattribute 5) from the RADIUS Framed-Route attribute [22].
- **$junos-framed-route-ipv6-distance**—Distance of an IPv6 access route; the variable is dynamically replaced with the preference value (Subattribute 5) from the RADIUS Framed-IPv6-Route attribute [99].

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Access Routes for Subscriber Management on page 690

preference (Tunnel Profile)

Syntax preference *number*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the preference for a tunnel. You can specify up to 8 levels of preference, and you can assign the same preference to a maximum of 31 tunnels. When you define multiple preferences for a destination, you increase the probability of a successful connection.

This value can be overridden by RADIUS attribute Tunnel-Preference [83].

Options *number*—Number that indicates the order in which the router attempts to connect to the destination. Zero is the highest level of preference.

Range: 0 through 2000

Default: 2000

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

 Copyright © 2013, Juniper Networks, Inc.

preferred-lifetime

Syntax preferred-lifetime *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix active. The DHCPv6 server sends this value to the client (router).

If a value for the **preferred-lifetime** statement is configured along with a value for the **maximum-lease-time** statement, or if the configured value of the **preferred-lifetime** statement is higher than the value of the **valid-lifetime** statement, a commit error occurs. If a value for the **preferred-lifetime** statement is not configured, Junos OS uses the same value as the **maximum-lease-time** or **valid-lifetime** statement. If neither the **maximum-lease-time** nor **valid-lifetime** statement is configured, Junos OS uses the default value of 86,400 seconds.

Options ***seconds***—Length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix 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 168
- maximum-lease-time on page 1778
- valid-lifetime on page 2102

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*

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 167
- Configuring Address-Assignment Pools on page 168
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*
- *Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)*

 Copyright © 2013, Juniper Networks, Inc.

prefix (DHCP Relay Agent)

Syntax prefix *prefix*;

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* forwarding-options dhcp-relay dhcpv6 group *group-name* relay-agent-interface-id],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* relay-option-82 circuit-id],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-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 group *group-name* relay-agent-interface-id],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-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 group *group-name* relay-option-82 circuit-id],
[edit 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 group *group-name* relay-agent-interface-id],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* relay-option-82 circuit-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.

Description Add a prefix to the base option 82 Agent Circuit ID information in DHCP packets destined for a DHCP server. The prefix can consist of any combination of the hostname, logical system name, and routing instance name. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

If you include only the hostname, only the logical system name, or only the routing instance name in the prefix, the format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces with stacked virtual LANs (S-VLANs) is one of the following:

host-name:(fe | ge)-*fpc/pic/port*:*svlan-id-vlan-id*
logical-system-name:(fe | ge)-*fpc/pic/port*:*svlan-id-vlan-id*
routing-instance-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, the format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs is as follows:

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, the format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs is as follows:

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 virtual LANs (VLANs) but not S-VLANs, only the ***vlan-id*** value appears in the Agent Circuit ID format. For Fast Ethernet or Gigabit Ethernet interfaces that do not use VLANs or S-VLANs, neither the ***vlan-id*** value nor the ***svlan-id*** value appears.

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 Agent Circuit ID information.
- **logical-system-name**—Prepend the name of the logical system to the Agent Circuit ID information.
- **routing-instance-name**—Prepend the name of the routing instance to the Agent Circuit ID information.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling and Disabling Insertion of Option 82 Information on page 325
- Configuring an Option 82 Prefix on page 326
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 329

Copyright © 2013, 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.

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*

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 ANCP 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 ANCP on page 1332
- Configuring ANCP for Backward Compatibility on page 1336
- Configuring ANCP Neighbors on page 1333

premium (Hierarchical Policer)

Syntax

```
premium {
  if-exceeding {
    bandwidth-limit bandwidth;
    burst-size-limit burst;
  }
  then {
    discard;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall hierarchical-policer],
[edit firewall hierarchical-policer]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... hierarchical-policer *name*]** hierarchy level introduced in Junos OS Release 11.4.

Description

On M40e, M120, and M320 edge routers with FPC input as FFPC and FPC output as SFPC, and on MX Series, T320, T640, and T1600 edge routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify a premium level for a hierarchical policer.

Options

Options are described separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Applying Policers*
- *Junos OS Class of Service Library for Routing Devices*
- *Hierarchical Policer Configuration Overview*
- *Hierarchical Policers*
- aggregate (Hierarchical Policer) on page 1447
- bandwidth-limit (Hierarchical Policer) on page 1475
- burst-size-limit (Hierarchical Policer) on page 1484
- hierarchical-policer on page 1675
- if-exceeding (Hierarchical Policer) on page 1692

Copyright © 2013, Juniper Networks, Inc.

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 lower 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 471
- Configuring Diameter Network Elements on page 473

priority (Dynamic Schedulers)

Syntax priority (*priority-level* | $junos-cos-scheduler-priority);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]

Release Information Statement introduced in Junos OS Release 9.3.
The **$junos-cos-scheduler-priority** predefined variable introduced in Junos OS Release 9.4.

Description Specify packet-scheduling priority value in a dynamic profile.

Options ***priority-level***—one of the following packet-scheduling priority values:

- **low**—Scheduler has low priority.
- **medium-low**—Scheduler has medium-low priority.
- **medium-high**—Scheduler has medium-high priority.
- **high**—Scheduler has high priority. Assigning high priority to a queue prevents the queue from being underserved.
- **strict-high**—Scheduler has strictly high priority. Configure a **high** priority queue with unlimited transmission bandwidth available to it. As long as it has traffic to send, the **strict-high** priority queue receives precedence over **low**, **medium-low**, and **medium-high** priority queues, but not **high** priority queues. You can configure **strict-high** priority on only one queue per interface.

$junos-cos-scheduler-priority—Junos predefined variable that is replaced with the packet-scheduling priority value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973
- Dynamic Variables Overview on page 640
- **scheduler (Dynamic Scheduler Maps) on page 1968**

 Copyright © 2013, 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.

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.

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 227
- Overriding Default DHCP Local Server Configuration Settings on page 220
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

 Copyright © 2013, Juniper Networks, Inc.

profile (Access)

Syntax

```
profile profile-name {
  accounting {
    address-change-immediate-update
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    coa-immediate-update;
    coa-no-override service-class-attribute;
    duplication;
    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;
    }
  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-per-tunnel number;
      multilink {
        drop-timeout milliseconds;
        fragment-threshold bytes;
      }
      ppp-authentication (chap | pap);
      ppp-profile profile-name;
      shared-secret shared-secret;
    }
    pap-password pap-password;
    ppp {
      cell-overhead;
      encapsulation-overhead bytes;
      framed-ip-address ip-address;
      framed-pool framed-pool;
      idle-timeout seconds;
```

```
        interface-id interface-id;
        keepalive seconds;
        primary-dns primary-dns;
        primary-wins primary-wins;
        secondary-dns secondary-dns;
        secondary-wins secondary-wins;
      }
      user-group-profile profile-name;
    }
    domain-name-server;
    domain-name-server-inet;
    domain-name-server-inet6;
    provisioning-order (gx-plus | jsrc);
    radius {
      accounting-server [ ip-address ];
      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;
        }
        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;
        }
        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 {
          agent-circuit-id;
          agent-remote-id;
          interface-description;
          nas-identifier;
```

 Copyright © 2013, Juniper Networks, Inc.

```
            }
            nas-port-type {
              ethernet {
                port-type;
              }
            }
            revert-interval interval;
            vlan-nas-port-stacked-format;
          }
          attributes {
            exclude {
              ...
            }
            ignore {
              framed-ip-netmask;
              input-filter;
              logical-system:routing-instance;
              output-filter;
            }
          }
        }
        radius-server server-address {
          accounting-port port-number;
          port port-number;
          retry attempts;
          routing-instance routing-instance-name;
          secret password;
          max-outstanding-requests value;
          source-address source-address;
          timeout seconds;
        }
        service {
          accounting-order (activation-protocol | radius);
        }
        session-options {
          client-group [ group-names ];
          client-idle-timeout minutes;
          client-session-timeoutminutes;
        }
      }
```

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.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- *Configuring the PPP Authentication Protocol*
- *Configuring Access Profiles for L2TP or PPP Parameters*
- *Configuring L2TP Properties for a Client-Specific Profile*
- Configuring an L2TP LNS with Inline Service Interfaces on page 417
- *Configuring PPP Properties for a Client-Specific Profile*
- Configuring Service Accounting with JSRC on page 494
- AAA Service Framework Overview on page 21
- *show network-access aaa statistics*
- *clear network-access aaa statistics*

promiscuous-mode (Protocols IGMP)

Syntax
promiscuous-mode;

Hierarchy Level
[edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.2 for dynamic profiles.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description
Specify that the interface accepts IGMP reports from hosts on any subnetwork. Note that when enabling promiscuous-mode, all routers on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring a Dynamic Profile for Client Access on page 679
- *Accepting IGMP Messages from Remote Subnetworks*

Copyright © 2013, Juniper Networks, Inc.

protocol (Dynamic Schedulers)

Syntax protocol (any | non-tcp | tcp);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name* drop-profile-map]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the protocol type for the specified scheduler in a dynamic profile.

Options **any**—Accept any protocol type.

non-tcp—Accept any protocol type other than TCP/IP.

tcp—Accept only TCP/IP protocol.

NOTE: Protocol types non-tcp and tcp are not supported on MX Series routers.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973

protocol (Subscriber Secure Policy)

Syntax protocol *protocol*;

Hierarchy Level [edit services radius-flow-tap policy *policy-name* inet drop-policy *rule-name* from],
[edit services radius-flow-tap policy *policy-name* inet6 drop-policy *rule-name* from]

Release Information Statement introduced in Junos OS Release 12.3.

Description Specify the match IP protocol type for the radius-flow-tap policy.

Options ***protocol***—Protocol for the IPv4 or IPv6 address for the radius-flow-tap policy.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

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 © 2013, 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 statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- For general information about configuring IGMP or MLD, see the *Multicast Protocols Feature Guide for Routing Devices*.

provisioning-order

Syntax provisioning-order (gx-plus | jsrc);

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.

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 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.

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 491
- Provisioning Subscribers with JSRC on page 494
- Configuring Gx-Plus on page 549
- Provisioning Subscribers with Gx-Plus on page 552

proxy-arp

Syntax proxy-arp;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.

Description For Ethernet interfaces only, configure the router to respond to any ARP request, as long as the router has an active route to the target address of the ARP request.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Restricted and Unrestricted Proxy ARP*
- *Configuring Gratuitous ARP*

 Copyright © 2013, 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.

The extended DHCP relay proxy is not supported for the J Series routers DHCP server. Also, 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 280
- Extended DHCP Relay Agent Overview on page 276
- Enabling DHCP Relay Proxy Mode on page 329

Copyright © 2013, Juniper Networks, Inc.

ps0 (Pseudowire Subscriber Interfaces)

Syntax

```
ps0 {
  anchor-point lt-device;
  mtu bytes;
  no-gratuitous-arp-request;
  (flexible-vlan-tagging | stacked-vlan-tagging | untagged | vlan-tagging);
}
```

Hierarchy Level

[edit logical-systems transport-ls interfaces]

Release Information

Statement introduced in Junos OS Release 13.1.

Description

Configure the pseudowire logical device.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Configuring a Pseudowire Subscriber Logical Interface Device on page 935
- Configuring the Transport Logical Interface for a Pseudowire Subscriber Logical Interface on page 937
- Configuring the Service Logical Interface for a Pseudowire Subscriber Logical Interface on page 939

Copyright © 2013, Juniper Networks, Inc.

pseudowire-service (Pseudowire Subscriber Interfaces)

Syntax

```
pseudowire-service {
    device-count number;
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure properties for the pseudowire devices on the router.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Pseudowire Subscriber Logical Interfaces Overview on page 923
- Configuring a Pseudowire Subscriber Logical Interface on page 933
- Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router on page 935

push (Dynamic VLANs)

Syntax push;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]

Release Information Statement introduced in Junos OS Release 10.4.

Description For dynamic VLAN interfaces, specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag. If you include the **push** statement in the configuration, you must also include the *pop* statement at the **[edit** dynamic-profiles ***profile-name*** interfaces ***interface-name*** unit ***logical-unit-number*** output-vlan-map**]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution*

qos-adjust

Syntax qos-adjust;

Hierarchy Level [edit protocols ancp]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify that the ANCP agent reports net data rates for downstream traffic to CoS.When this statement is not configured, ANCP does not report traffic rates to CoS.

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 1338
- Traffic Rate Reporting and Adjustment by ANCP on page 1314
- Configuring ANCP on page 1332

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 ANCP-reported downstream and upstream data rate 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 1337
- Configuring ANCP on page 1332

 Copyright © 2013, 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 ANCP-reported downstream and upstream data rate 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 1337
- Configuring ANCP on page 1332

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 ANCP-reported downstream and upstream data rate 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 1337
- Configuring ANCP on page 1332

Copyright © 2013, Juniper Networks, Inc.

qos-adjust-sds1

Syntax qos-adjust-sds1 *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 ANCP-reported downstream and upstream data rate for all subscribers on an SDS1 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 1337
- Configuring ANCP on page 1332

qos-adjust-vds1

Syntax qos-adjust-vds1 *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 ANCP-reported downstream and upstream data rate for all subscribers on an VDS1 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 1337
- Configuring ANCP on page 1332

 Copyright © 2013, Juniper Networks, Inc.

qos-adjust-vds2

Syntax qos-adjust-vds2 *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 ANCP-reported downstream and upstream data rate for all subscribers on an VDS2 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 1337
- Configuring ANCP on page 1332

qualified-bum-pruning-mode

Syntax qualified-bum-pruning-mode;

Hierarchy Level [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 10.4.

Description For Junos OS Layer 2 Wholesale configurations, prune (constrain) distribution of broadcast, unicast, and multicast (BUM) packets of unknown origin to only those interfaces that match the traffic from a specific VLAN pair.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers*
- *Junos OS VPNs Library for Routing Devices*

qualified-next-hop (Subscriber Management)

Syntax

```
qualified-next-hop interface-name {
    mac-address address;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access-internal route *subscriber-ip-address*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access-internal route *subscriber-ip-address*],
[edit dynamic-profiles *profile-name* routing-options access-internal route *subscriber-ip-address*]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options route *subscriber-ip-address*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* route *subscriber-ip-address*]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure the qualified next-hop and the MAC address for an access-internal route for DHCP and PPP subscriber interfaces.

Options

interface-name—Either the specific interface you want to assign to the access route or the variable, or the **$junos-interface-name** variable. The variable is dynamically replaced with the value supplied by DHCP or PPP when a subscriber logs in.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 692

Copyright © 2013, 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;
    }
    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 {
      agent-circuit-id;
```

```
            agent-remote-id;
            interface-description;
            nas-identifier;
        }
        nas-port-type {
            ethernet {
                port-type;
            }
        }
        revert-interval interval;
        vlan-nas-port-stacked-format;
    }
}
```

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.

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 36
- RADIUS Server Options for Subscriber Access on page 41

 Copyright © 2013, Juniper Networks, Inc.

radius (Dynamic Profiles)

Syntax

```
radius {
  vendor-id id {
    attribute attribute-number;
    tag tag-number;
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* variables]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure RADIUS attribute variables in a dynamic profile.

The statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring User-Defined CoS Variables in a Dynamic Service Profile on page 1000

radius-coa (Adjustment Control Profiles)

Syntax

```
radius-coa {
   priority priority;
   algorithm algorithm;
}
```

Hierarchy Level

[edit class-of-service adjustment-control-profiles *profile-name* application]

Release Information

Statement introduced in Junos OS Release 13.1.

Description

Configure the shaping rate adjustment controls for the RADIUS CoA application.

Options

priority—Priority of the RADIUS CoA application in the adjustment control profile.

Range: 1 through 10; 1 being the highest priority.

Default: 1

algorithm—Rate adjustment algorithm used by the RADIUS CoA application.

Values:

- adjust-never—Do not perform rate adjustments.
- adjust-always—Adjust the shaping rate unconditionally.
- adjust-less—Adjust the shaping rate if it is less than the configured value.
- adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
- adjust-greater—Adjust the shaping rate if it is greater than the configured value.
- adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.

Default: adjust-always

Required Privilege Level

interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- CoS Adjustment Control Profiles Overview on page 1084
- Configuring CoS Adjustment Control Profiles on page 1106
- Verifying the CoS Adjustment Control Profile Configuration on page 1107
- **adjustment-control-profiles on page 1442**
- **application (Adjustment Control Profiles) on page 1459**

Copyright © 2013, 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 Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 247

Copyright © 2013, 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;
        }
      }
    }
  }
  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.

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 on page 1245
- Configuring Support for Subscriber Secure Policy Mirroring on page 1264

 Copyright © 2013, Juniper Networks, Inc.

radius-options (Edit Access)

Syntax

```
radius-options {
  revert-interval seconds;
  request-rate rate;
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 8.3.

Description Configure RADIUS options.

The remaining statements are explained separately.

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 47
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24

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.

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 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65

Copyright © 2013, Juniper Networks, Inc.

radius-server

Syntax

```
radius-server server-address {
   accounting-port port-number;
   port port-number;
   retry attempts;
   routing-instance routing-instance-name;
   secret password;
   max-outstanding-requests value;
   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.

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—Address of the RADIUS authentication server.

The remaining statements are explained separately.

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 RADIUS Authentication*
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- *Configuring an EX Series Switch to Use Junos Pulse Access Control Service for Network Access Control (CLI Procedure)*
- *show network-access aaa statistics*
- *clear network-access aaa statistics*

Copyright © 2013, Juniper Networks, Inc.

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 167
- Configuring Address-Assignment Pools on page 168
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

Copyright © 2013, 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 228
- Overriding Default DHCP Local Server Configuration Settings on page 220

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*

rebalance-periodic (Aggregated Ethernet Subscriber Interfaces)

Syntax rebalance-periodic time *hour:minute* <interval *hours*>

Hierarchy Level [edit interfaces ae *number* aggregated-ether-options]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure periodic rebalancing of distribution of subscribers on an aggregated Ethernet bundle.

Options ***hour:minute***—Time at which the rebalancing occurs, in military time.

hours—Interval at which the rebalancing occurs, in hours. Default: 24 hours.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface on page 830

 Copyright © 2013, Juniper Networks, Inc.

realm (Diameter Origin)

Syntax realm *realm-name*;

Hierarchy Level [edit diameter origin]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the realm of the host that originates the Diameter message.

Options ***realm-name***—Name of the message origin realm. Supplied as the value of Origin-Realm AVP for all messages sent by 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 471
- Configuring the Origin Attributes of the Diameter Instance on page 472

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.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

registration-lifetime

Syntax

registration-lifetime *seconds*;

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit services mobile-ip home-agent virtual-network home-agent-address *ip-address*]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Configure maximum period for registration lifetime that is accepted by the Mobile IP home agent.

Options

registration-lifetime *seconds*—Maximum lifetime that the home agent accepts in any registration request. The registration lifetime is not affected if you change the system clock.

Range: 7 through 65535 seconds

Default: 3600 seconds

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

relay-agent-interface-id (DHCPv6 Relay Agent)

Syntax

```
relay-agent-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 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.

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.

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 1548
- Extended DHCP Relay Agent Overview on page 276
- DHCPv6 Relay Agent Overview on page 279
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 329

Copyright © 2013, 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 279
- Creating Unique Usernames for DHCP Clients on page 239

relay-agent-remote-id (DHCPv6 Relay Agent)

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 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.

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.

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 279
- Creating Unique Usernames for DHCP Clients on page 239

Copyright © 2013, Juniper Networks, Inc.

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 239

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.

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.

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 239

Copyright © 2013, 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 239

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 279
- Creating Unique Usernames for DHCP Clients on page 239

 Copyright © 2013, 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.

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 323

relay-option-82

Syntax

```
relay-option-82 {
  circuit-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
}
```

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.

If you enable insertion of option 82 information in DHCP packets, you must specify at least the **circuit-id** statement to include the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option.

You can use the **relay-option-82** statement and its subordinate statements at the **[edit forwarding-options dhcp-relay]** hierarchy level to control insertion of option 82 information globally, or at the **[edit forwarding-options dhcp-relay group *group-name*]** hierarchy level to control insertion of option 82 information for a named group of interfaces.

To restore the default behavior (option 82 information is not inserted into DHCP packets), use the **delete relay-option-82** statement.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling and Disabling Insertion of Option 82 Information on page 325
- dhcp-relay on page 1548

Copyright © 2013, Juniper Networks, Inc.

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 323

remote-gateway (Tunnel Profile)

Syntax

```
remote-gateway {
  address server-ip-address;
  gateway-name server-name;
}
```

Hierarchy Level

[edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Specify the IP address and hostname of the remote gateway at the L2TP tunnel endpoint, the LNS.

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

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]

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 168

Copyright © 2013, 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 276
- Replacing the DHCP Relay Request and Release Packet Source Address on page 295

replay-method

Syntax

replay-method (none | timestamp *seconds*);

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*].
[edit logical-systems *logical-system-name* services mobile-ip peer *nai*@*domain* spi *hexadecimal-value*].
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*].
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer *nai*@*domain* spi *hexadecimal-value*].
[edit routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*].
[edit routing-instances *routing-instances-name* services mobile-ip peer *nai*@*domain* spi *hexadecimal-value*],
[edit services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit services mobile-ip peer *nai*@*domain* spi *hexadecimal-value*]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Configure the replay protection method. The Identification field enables the home agent to verify that a registration message has been recently generated by the mobile node, rather than replayed by an attacker from a previous registration. You can specify a timestamp tolerance for the mobile node, which causes the request to be rejected if the tolerance is exceeded, or you can specify that the tolerance be taken from the value configured on the home agent.

Default

If you do not configure the replay protection method, then the timestamp tolerance is taken from the home agent by default.

Options

none—Timestamp tolerance is obtained from the setting configured for the home agent

timestamp *seconds*—Tolerance time in which a registration request timestamp and the local time of the home agent can differ.

Range: 1 through 255 seconds

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

Copyright © 2013, 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 1625
- Configuring RADIUS Server Options for Subscriber Access on page 47
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- Subscriber Access Interface Description Storage and Reporting Through RADIUS Overview on page 50

request-rate

Syntax request-rate *rate*;

Hierarchy Level [edit access radius-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description (M120, M320, and MX Series routers) 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: 500 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

- Configuring RADIUS Server Options for Subscriber Access on page 47
- Configuring Router or Switch Interaction with RADIUS Servers on page 23

retransmission-count-established (L2TP)

Syntax retransmission-count-established *count*;

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 12.1.

Description Set the maximum number of times control messages are retransmitted for established tunnels.

BEST PRACTICE: **Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing no services l2tp tunnel retransmission-count-established.**

Options *count*—Number of retransmissions.

Range: 2 through 30

Default: 7

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Number of L2TP Control Message Retransmissions on page 411
- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

 Copyright © 2013, Juniper Networks, Inc.

retransmission-count-not-established (L2TP)

Syntax retransmission-count-not-established *count*;

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 12.1.

Description Set the maximum number of times control messages are retransmitted for tunnels that are not established.

BEST PRACTICE: **Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing no services l2tp tunnel retransmission-count-not-established.**

Options *count*—Number or retransmissions.

Range: 2 through 30

Default: 5

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Number of L2TP Control Message Retransmissions on page 411
- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

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*

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.

Options ***attempts***—Number of times that the router is allowed to attempt to contact a RADIUS server.

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 Authentication and Accounting Parameters for Subscriber Access on page 24
- Configuring Router or Switch Interaction with RADIUS Servers on page 23
- *Example: Configuring CHAP Authentication with RADIUS*
- *Configuring RADIUS Authentication for L2TP*
- **timeout on page 2025**

 Copyright © 2013, 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 604800 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 47
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24

revocation-required

Syntax revocation-required;

Hierarchy Level [edit logical-systems *logical-system-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent virtual-networkhome-agent-address *ip-address*],
[edit services mobile-ip home-agent virtual-network home-agent-address *ip-address*]

Release Information Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Configure the Mobile IP home agent to accept registration revocation requests only when the request includes the revocation extension.

Default The Mobile IP home agent supports registration revocation requests that include the revocation extension, but it does not require the extension.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

 Copyright © 2013, Juniper Networks, Inc.

rewrite-rules (Dynamic CoS Interfaces)

Syntax

```
rewrite-rules {
  dscp (rewrite-name | default);
  dscp-ipv6 (rewrite-name | default);
  ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
  inet-precedence (rewrite-name | default);
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 10.1.

Description

Associate a rewrite-rules configuration or default mapping with a specific interface in a dynamic profile.

Options

rewrite-name—Name of a **rewrite-rules** mapping configured at the **[edit class-of-service rewrite-rules]** hierarchy level.

default—The default mapping.

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

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- *rewrite-rules*

route (Access)

Syntax

```
route prefix {
  next-hop next-hop;
  metric route-cost;
  preference route-distance;
  tag route-tag;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access],
[edit dynamic-profiles *profile-name* routing-options access]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure the parameters for access routes.

Options

prefix—Either the specific route prefix that you want to assign to the access route or one of the following route prefix variables.

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

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 690

Copyright © 2013, Juniper Networks, Inc.

route (Access Internal)

Syntax

```
route subscriber-ip-address {
  next-hop next-hop;
  qualified-next-hop underlying-interface {
    mac-address address;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access-internal],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access-internal],
[edit dynamic-profiles *profile-name* routing-options access-internal]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access-internal]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access-internal]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure parameters for an access-internal route.

Options

subscriber-ip-address—Either the specific IP address you want to assign to the access-internal route or the subscriber IP address variable ($junos-subscriber-ip-address). The subscriber IP address variable is dynamically replaced with the value supplied by DHCP or PPP when a subscriber logs in.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 692
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 692

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.

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 471
- Configuring Diameter Network Elements on page 473

Copyright © 2013, Juniper Networks, Inc.

route-suppression (DHCP)

Syntax

route-suppression (access | access-internal | destination);

Hierarchy Level

[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group *group-name*],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name*],
[edit logical-systems *logical-system-name* ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*...],
[edit routing-instances *routing-instance-name* ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server group *group-name*],
[edit system services dhcp-local-server dhcpv6 group *group-name*]

Release Information

Statement introduced in Junos OS Release 13.2.

Description

Configure the jdhcpd process to suppress the installation of access, access-internal, or destination routes during client binding.

Options

access—(DHCPv6 only) Suppress installation of access routes. You can use the **access** and **access-internal** options in the same statement for DHCPv6.

access-internal—In a DHCPv4 hierarchy, suppress installation of both access-internal and destination routes. In a DHCPv6 hierarchy, suppress access-internal routes only. Can be configured in the same statement with the **access** option.

destination—(DHCPv4 only) Suppress installation of destination routes. This option and the **access-internal** option are mutually exclusive; however, the **access-internal** option also suppresses destination routes.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default on page 223

router (Address-Assignment Pools)

Syntax router [*router-address*];

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes]

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 168

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.

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 © 2013, Juniper Networks, Inc.

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 24

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 471
- Configuring Diameter Peers on page 473

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 474

routing-instance (Tunnel Profile)

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify a routing instance for a tunnel.

Options ***routing-instance-name***—Name of the routing instance.
Default: Routing instance *default*

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 403

 Copyright © 2013, Juniper Networks, Inc.

routing-instance (PPPoE Service Name Tables)

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Use in conjunction with the **dynamic-profile** statement at the same hierarchy levels to specify the routing instance in which to instantiate a dynamic PPPoE interface. You can associate a routing instance with a named service entry, **empty** service entry, or **any** service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The routing instance associated with a service entry in a PPPoE service name table overrides the routing instance associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the **routing-instance** statement at the **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level, you cannot also include the **static-interface** statement at this level. The **routing-instance** and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options ***routing-instance-name***—Name of the routing instance in which the router instantiates the dynamic PPPoE interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring PPPoE Service Name Tables*
- Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation on page 910

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 © 2013, Juniper Networks, Inc.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214

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 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.

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.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214
- Creating Unique Usernames for DHCP Clients on page 239

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 500
- Configuring the Static Subscriber Global Username on page 503
- Configuring the Static Subscriber Group Username on page 507

Copyright © 2013, Juniper Networks, Inc.

routing-instances

Syntax

```
routing-instances routing-instance-name {
  interface interface-name;
  routing-options {
    access {
      route prefix {
        metric route-cost;
        next-hop next-hop;
        preference route-distance;
        tag route-tag;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
    multicast {
      interface interface-name {
        no-qos-adjust;
      }
    }
    rib routing-table-name {
      access {
        route prefix {
          metric route-cost;
          next-hop next-hop;
          preference route-distance;
          tag route-tag;
        }
      }
      access-internal {
        route subscriber-ip-address {
          qualified-next-hop underlying-interface {
            mac-address address;
          }
        }
      }
    }
  }
}
```

Hierarchy Level [edit dynamic-profiles]

Release Information Statement introduced in Junos OS Release 9.6.
The **routing-options** statement introduced in Junos OS Release 10.1.

Description Dynamically configure an additional routing entity for a router.

Options ***routing-instance-name***—The routing instance variable (*$junos-routing-instance*). The routing instance variable is dynamically replaced with the routing instance the accessing client uses when connecting to the router.

NOTE: **Though we do not recommend it, you can also enter a specific name for the routing instance, a maximum of 31 characters.**

The remaining statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution*

 Copyright © 2013, Juniper Networks, Inc.

routing-options (Dynamic Profiles)

Syntax

```
routing-options {
  access {
    route prefix {
      metric route-cost;
      next-hop next-hop;
      preference route-distance;
      tag route-tag;
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
  multicast {
    interface interface-name {
      no-qos-adjust;
    }
  }
  rib routing-table-name {
    access {
      route prefix {
        metric route-cost;
        next-hop next-hop;
        preference route-distance;
        tag route-tag;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance]

Release Information

Statement introduced in Junos OS Release 9.6.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance]** hierarchy level introduced in Junos OS Release 10.1.

Description

Configure protocol-independent routing properties in a dynamic profile.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Access Routes for Subscriber Management on page 690
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 692

rpf-check (Dynamic Profiles)

Syntax

```
rpf-check {
   fail-filter filter-name;
   mode loose;
}
```

Hierarchy Level
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information
Statement introduced in Junos OS Release 9.6.

Description
Check whether traffic is arriving on an expected path. You can include this statement with the **inet** protocol family 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
- *Configuring Unicast RPF*
- Configuring Unicast RPF and Fail Filters in Dynamic Profiles for Subscriber Interfaces on page 1179

Copyright © 2013, Juniper Networks, Inc.

rule

Syntax

```
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      application [junos-http, junos-https, junos-httpproxy];
      destination-address address <except>;
      destination-prefix-list list-name <except>;
    }
    then {
      accept;
      rewrite {
        destination-address address;
        destination-port port;
      }
      syslog;
    }
  }
}
```

Hierarchy Level

[edit services captive-portal-content-delivery]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Specify the rule the router uses when applying this service.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately.

Required Privilege Level

services—To view this statement in the configuration.
services–control—To add this statement to the configuration.

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

rule (IP Reassembly)

Syntax

```
rule rule-name {
   match-direction direction;
}
```

Hierarchy Level

[edit services ip-reassembly]

Release Information

Statement introduced in Junos OS Release 13.1.

Description

Configure an IP reassembly rule, which is used for inline IP reassembly on the inline services interface. The IP reassembly rule can be attached to a service set to indicate that the service set is of type IP reassembly. For inline IP reassembly, each rule must include the **match-direction** statement, which specifies the direction in which the match is applied.

NOTE: **If you configure an IP reassembly rule, then you must configure the match-direction statement.**

Options

rule-name—Name of the IP reassembly rule.

Range: Up to 63 characters

The remaining statement is explained separately.

NOTE: **To create more than one IP reassembly rule, include the rule statement multiple times.**

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring IP Inline Reassembly for L2TP on page 421
- IP Packet Fragment Reassembly for L2TP Overview on page 390

Copyright © 2013, Juniper Networks, Inc.

rule-set

Syntax

```
rule-set rule-set-name {
   [rule rule-name ];
}
```

Hierarchy Level [edit services (captive-portal-content-delivery) captive-portal-content-delivery]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the rule set the router uses when applying this service.

Options ***rule-set-name***—Identifier for the collection of rules that constitute this rule set.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

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

Syntax rx-connect-speed-when-equal

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 13.1.

Description Enable sending the receive connect speed, which is represented by AVP 38, even when its value is equal to that of the transmit connect speed, which is represented by AVP 24. By default, AVP 38 is sent from the LAC to the LNS during the establishment of an L2TP tunnel session, only when its value is different from AVP 24. You can override the default setting with this configuration statement.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method to Set the LAC Connection Speeds to the LNS on page 408
- Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds are Equal on page 409

scheduler (Dynamic Scheduler Maps)

Syntax scheduler *scheduler-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service scheduler-maps *map-name* forwarding-class *class-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Associate a scheduler with a scheduler map in a dynamic profile.

Options *scheduler-name*—Either the specific name of the scheduler configuration block or the scheduler variable (**$junos-cos-scheduler**).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973
- Dynamic Variables Overview on page 640

scheduler-map (Dynamic Traffic Shaping)

Syntax scheduler-map (*map-name*);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.3.
The **$junos-cos-scheduler-map** variable introduced in Junos OS Release 9.4.

Description Associate a scheduler map name with a traffic-control profile in a dynamic profile.

The scheduler map can be defined dynamically (at the **[edit dynamic-profiles *profile-name* class-of-service scheduler-maps]** hierarchy level) or statically (at the **[edit class-of-service scheduler-maps]** hierarchy level).

Options *map-name*—Name of the scheduler map or the Junos predefined variable (**$junos-cos-scheduler-map**). When you specify the variable, the scheduler-map name is obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Traffic Scheduling and Shaping for Subscriber Access on page 971
- output-traffic-control-profile on page 1839

 Copyright © 2013, Juniper Networks, Inc.

scheduler-maps (Dynamic CoS Definition)

Syntax

```
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the [**edit dynamic-profiles *profile-name***] hierarchy level introduced in Junos OS Release 9.3.

Description

Specify a scheduler map name in a dynamic profile and associate it with the scheduler configuration and forwarding class.

Options

map-name—Name of the scheduler map.

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

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973

Copyright © 2013, Juniper Networks, Inc.

schedulers (Dynamic CoS Definition)

Syntax

```
schedulers {
  scheduler-name{
    adjust-minimum rate;
    adjust-percent percentage;
    buffer-size (percent percentage | remainder | temporal microseconds |
      $junos-cos-scheduler-bs);
    drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol
      (any | non-tcp | tcp) drop-profile (profile-name | predefined-variable);
    excess-priority (low | high | $junos-cos-scheduler-excess-priority | none);
    excess-rate (percent percentage  | percent $junos-cos-scheduler-excess-rate);
    priority (priority-level | $junos-cos-scheduler-priority);
    shaping-rate (rate | predefined-variable) <burst-size  bytes>;
    transmit-rate (rate | percent percentage | remainder | percent percentage
      $junos-cos-scheduler-tx) <exact | rate-limit>;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service]

Release Information

Statement introduced in Junos OS Release 9.3.
The **$junos-cos-scheduler** predefined variable introduced in Junos OS Release 9.4.

Description

Specify scheduler name and parameter values in a dynamic profile.

Options

scheduler-name—Name of the scheduler to be configured or the Junos OS predefined variable (**$junos-cos-scheduler**). The predefined variable is replaced with the scheduler name obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973
- **scheduler on page 1968**

Copyright © 2013, 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 24
- Configuring Router or Switch Interaction with RADIUS Servers on page 23
- *Example: Configuring CHAP Authentication with RADIUS*
- *Configuring RADIUS Authentication for L2TP*
- *Configuring the RADIUS Disconnect Server for L2TP*
- *Configuring an EX Series Switch to Use Junos Pulse Access Control Service for Network Access Control (CLI Procedure)*

secret (Tunnel Profile)

Syntax secret *password*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the tunnel password sent to the LNS for authentication.

Options ***password***—Cleartext password.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

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.

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 36
- Configuring Per-Subscriber Session Accounting on page 29

 Copyright © 2013, 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 276
- Overriding the Default DHCP Relay Configuration Settings on page 293
- Sending Release Messages When Clients Are Deleted on page 312

server (Dynamic PPPoE)

Syntax server;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" pppoe-options]

Release Information Statement introduced in Junos OS Release 10.1.

Description In a dynamic profile, configure the router to act as a PPPoE server, also known as a remote access concentrator, when a PPPoE logical interface is dynamically created.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Basic PPPoE Dynamic Profile on page 900
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

 Copyright © 2013, 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 1548
- Extended DHCP Relay Agent Overview on page 276
- Configuring Server Groups on page 328

server-identifier (Address-Assignment Pools)

Syntax server-identifier *ipv4-address*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes]

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 168

service (Service Accounting)

Syntax

```
service {
   accounting-order (activation-protocol | radius);
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Define the subscriber service accounting configuration.

The remaining statement is explained separately.

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 494
- Service Accounting with JSRC on page 486

 Copyright © 2013, Juniper Networks, Inc.

service (Dynamic Service Sets)

Syntax

```
service {
  input {
    service-set service-set-name {
      service-filter filter-name;
    }
    post-service-filter filter-name;
  }
  output {
    service-set service-set-name {
      service-filter filter-name;
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos–interface–unit" family *family*]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family*]** hierarchy level introduced in Junos OS Release 10.1.

Description

Define the service sets and filters to be applied to an interface. This statement is not supported for **family inet6**.

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

- Dynamic Service Sets Overview on page 1151
- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197

service-device-pool (L2TP)

Syntax service-device-pool *pool-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Assign a pool of service interfaces to the tunnel group to balance traffic across.

NOTE: The service interface configuration is required for static LNS sessions. Either the service interface configuration or the service device pool configuration can be used for dynamic LNS sessions.

Options *pool-name*—Name of the service device pool.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429

service-device-pools (L2TP Service Interfaces)

Syntax

```
service-device-pools {
  pool pool-name {
    interface service-interface-name;
  }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure one or more pools of service interfaces that can be assigned to an L2TP tunnel group for traffic load-balancing. The service device pool is required for dynamic LNS sessions.

Options *pool-name*—Name of the service interface pool.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 430

 Copyright © 2013, Juniper Networks, Inc.

service-filter (Dynamic Service Sets)

Syntax

service-filter *filter-name*;

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* service input service-set *service-set-name*],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* service output service-set *service-set-name*],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* service input service-set *service-set-name*],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* service output service-set *service-set-name*]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* service input service-set *service-set-name*]** and **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* service output service-set *service-set-name*]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Define the filter to be applied to traffic before it is accepted for service processing. Configuration of a service filter is optional; if you include the **service-set** statement without a **service-filter** definition, the router software assumes that the match condition is true and selects the service set for processing automatically. Only the Internet Protocol version 4 (IPv4) protocol family is currently supported for dynamic PPPoE logical interfaces.

Options

filter-name—Identifies the filter to be applied in service processing.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Service Sets Overview on page 1151
- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197

Copyright © 2013, Juniper Networks, Inc.

service-interface

Syntax service-interface *interface-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.
Option **si-*fpc/pic/port*** introduced in Junos OS Release 11.4.

Description Specify the service interface responsible for handling L2TP processing.

NOTE: On MX Series routers, the service interface configuration is required for static LNS sessions. Either the service interface configuration or the service device pool configuration can be used for dynamic LNS sessions.

Options ***interface-name***—Name of the service interface. The interface type depends on the line card as follows:

- **sp-*fpc/pic/port***—On AS or Multiservices PICs on M7i, M10i, and M120 routers.
- **si-*fpc/pic/port***—On MPCs on MX Series routers.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Local Gateway Address and PIC*
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429

 Copyright © 2013, Juniper Networks, Inc.

service-name-table

Syntax service-name-table *table-name*;

Hierarchy Level
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information
Statement introduced in Junos OS Release 10.0.
Support at the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.

Description Specify the PPPoE service name table assigned to a PPPoE underlying interface. This underlying interface is configured with either the **encapsulation ppp-over-ether** statement or the **family pppoe** statement; the two statements are mutually exclusive.

NOTE: The **[edit ... family pppoe]** hierarchies are supported only on MX Series routers with MPCs.

Options ***table-name***—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring PPPoE Service Name Tables*
- *Assigning a Service Name Table to a PPPoE Underlying Interface*
- Configuring the PPPoE Family for an Underlying Interface on page 831

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 200
- Default Subscriber Service Overview on page 1127
- Configuring a Default Subscriber Service on page 1127

Copyright © 2013, 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 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.

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 1548
- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236
- Grouping Interfaces with Common DHCP Configurations on page 217
- Default Subscriber Service Overview on page 1127
- Configuring a Default Subscriber Service on page 1127

Copyright © 2013, Juniper Networks, Inc.

service-set (Dynamic Service Sets)

Syntax

```
service-set service-set-name {
    service-filter filter-name;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
    family service input],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
    family service output],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family
    service input],
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family family
    service output]
```

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* service input]** and **[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family *family* service output]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Define one or more service sets in a dynamic profile. Service sets are applied to an interface. If you define multiple service sets, the router software evaluates the filters in the order in which they appear in the configuration. Only the Internet Protocol version 4 (IPv4) protocol family is currently supported for dynamic PPPoE logical interfaces.

Options

service-set-name—Name of the service set.

The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Dynamic Service Sets Overview on page 1151
- Associating Service Sets with Interfaces in a Dynamic Profile on page 1197

Copyright © 2013, Juniper Networks, Inc.

services (captive-portal-content-delivery)

Syntax services captive-portal-content-delivery { ... }

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 10.4.

Description Define the captive portal and content delivery set of the rules statements to be applied to traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

session-mode

Syntax session-mode (automatic | multihop | singlehop);

Hierarchy Level [edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the session mode.

Options **automatic**—Configure single-hop BFD sessions if the peer is directly connected to the router interface and multihop BFD sessions if the peer is not directly connected to the router interface.

multihop—Configure multihop BFD sessions.

single-hop—Configure single hop BFD sessions.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

 Copyright © 2013, Juniper Networks, Inc.

session-options

Syntax

```
session-options {
   client-group [ group-names ];
   client-idle-timeout minutes;
   client-session-timeoutminutes;
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description (J Series, MX Series, and SRX Series devices) Define options that control a user's session after successful authentication.

The remaining statements are explained separately.

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 153
- Configuring Subscriber Session Options on page 154

shaping-rate (Dynamic Traffic Shaping and Scheduling)

Syntax shaping-rate (*rate* | *predefined-variable*) <burst-size *bytes* | $junos-cos-shaping-rate-burst>;

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*], [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]

Release Information Statement introduced in Junos OS Release 9.2.
The **$junos-cos-shaping-rate** variable for traffic-control profiles introduced in Junos OS Release 9.4.
The **$junos-cos-scheduler-shaping-rate** variable for schedulers introduced in Junos OS Release 10.2.
Option **burst-size** introduced in Junos OS Release 11.4.

Description Configure a shaping rate for a logical interface or a scheduler. The sum of the shaping rates for all logical interfaces on the physical interface can exceed the physical interface bandwidth. This practice is known as oversubscription of the peak information rate (PIR).

Options *rate*—Peak rate in bits per second (bps). You can specify the value as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1000 through 160,000,000,000 bps

predefined-variable—One of the following Junos predefined variables. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

- **$junos-cos-shaping-rate**—Variable for the shaping rate that is specified for the logical interface. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
- **$junos-cos-scheduler-shaping-rate**—Variable for the shaping rate that is specified for a scheduler. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]** hierarchy level.

burst-size *bytes*—(Optional) Maximum burst size, in bytes.

Range: 0 through 1,000,000,000

$junos-cos-shaping-rate-burst—(Optional) Variable for the burst-size that is specified for the shaping rate. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profile]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Traffic Scheduling and Shaping for Subscriber Access on page 971
- **output-traffic-control-profile on page 1839**

 Copyright © 2013, Juniper Networks, Inc.

shared-name

Syntax shared-name *filter-shared-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family-name* filter [input | output] *filter-name*]

Release Information Statement introduced in Junos OS Release 12.2.

Description Apply a filter shared name to a dynamic filter.

Options ***filter-shared-name***— Name of the specific shared filter or $junos-interface-set-name.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- For general information about configuring firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Library for Routing Devices*
- Dynamic Firewall Filters Overview on page 1132
- Classic Filters Overview on page 1133
- Basic Classic Filter Syntax on page 1135

shared-secret

Syntax shared-secret *shared-secret*;

Hierarchy Level [edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced in Junos OS Release 7.4.

Description Configure the shared secret.

Options ***shared-secret***—Shared secret key for authenticating the peer.

Required Privilege Level admin—To view this statement in the configuration.
admin–control—To add this statement to the configuration.

Related Documentation

- *Configuring L2TP Properties for a Client-Specific Profile*
- Configuring an L2TP Access Profile on the LNS on page 422

short-cycle-protection (Static and Dynamic Subscribers)

Syntax

short-cycle-protection <lockout-time-min *minimum-seconds* lockout-time-max *maximum-seconds*>;

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit interfaces demux0 unit *logical-unit-number* family pppoe]
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Configure the router to temporarily prevent (lock out) a failed or short-lived (also known as short-cycle) PPPoE subscriber session from reconnecting for a default or configurable period of time. You can optionally override the default lockout time, 1 through 300 seconds (5 minutes), by specifying the minimum lockout time and maximum lockout time as part of the **short-cycle-protection** statement.

Options

lockout-time-min *minimum-seconds*—(Optional) Minimum lockout time for failed or short-lived PPPoE subscriber sessions. The *minimum-seconds* value must be less than or equal to the *maximum-seconds* value. Setting *minimum-seconds* and *maximum-seconds* to the same value causes the lockout time to become fixed at that value.

Range: 1 through 86400 (24 hours)

Default: 1

lockout-time-max *maximum-seconds*—(Optional) Maximum lockout time for failed or short-lived PPPoE subscriber sessions. The *maximum-seconds* value must be equal to or greater than the *minimum-seconds* value. Setting *maximum-seconds* and *minimum-seconds* to the same value causes the lockout time to become fixed at that value.

Range: 1 through 86400 (24 hours)

Default: 300 (5 minutes)

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Lockout of PPPoE Subscriber Sessions on page 912
- PPPoE Subscriber Session Lockout Overview on page 893
- Understanding the Lockout Period for PPPoE Subscriber Session Lockout on page 896

Copyright © 2013, Juniper Networks, Inc.

- For information about configuring dynamic PPPoE subscriber interfaces, see the *Junos OS Subscriber Management and Services Library*
- For information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

single-rate

Syntax

```
single-rate {
  (color-aware | color-blind);
  committed-information-rate bps;
  committed-burst-size bytes;
  excess-burst-size bytes;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall three-color-policer *name*],
[edit firewall three-color-policer *policer-name*],
[edit logical-systems *logical-system-name* firewall three-color-policer *policer-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... three-color-policer *name*]** hierarchy level introduced in Junos OS Release 11.4.

Description

Configure a single-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).

Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the EBS are assigned medium-high loss priority (yellow). Packets that exceed the EBS are assigned high loss priority (red).

Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Three-Color Policer Configuration Overview*
- color-aware on page 1514
- color-blind on page 1515
- two-rate on page 2067

sip-server-address

Syntax sip-server-address *ipv6-address*;

Hierarchy Level [edit access address-assignment pool *pool-name* family *family* dhcp-attributes]

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 167
- Configuring Address-Assignment Pools on page 168

sip-server-domain-name

Syntax sip-server-domain-name *domain-name*;

Hierarchy Level [edit access address-assignment pool *pool-name* family *family* dhcp-attributes]

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 167
- Configuring Address-Assignment Pools on page 168

 Copyright © 2013, Juniper Networks, Inc.

source (Dynamic IGMP Interface)

Syntax source *source*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name* static]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the IP version 4 (IPv4) unicast address to send data on an interface.

Options ***source***—IPv4 unicast address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring a Dynamic Profile for Client Access on page 679
- For information about defining an IGMP source, see "Enabling IGMP Static Group Membership" in the *Multicast Protocols Feature Guide for Routing Devices*

source (Dynamic MLD Interface)

Syntax
```
source ip-address {
   source-count number;
   source-increment increment;
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 10.1.

Description IP version 6 (IPv6) unicast source address for the multicast group being configured on a dynamic interface.

Options ***ip-address***—One or more IPv6 unicast addresses.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Enabling MLD Static Group Membership*

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.

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 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 23
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- *Example: Configuring CHAP Authentication with RADIUS*
- *Configuring RADIUS Authentication for L2TP*

source-address (Subscriber Secure Policy)

Syntax source-address *address*;

Hierarchy Level [edit services radius-flow-tap policy *policy-name* inet drop-policy *rule-name* from],
[edit services radius-flow-tap policy *policy-name* inet6 drop-policy *rule-name* from]

Release Information Statement introduced in Junos OS Release 12.3.

Description Specify source IP address or prefix value from which to inherit configuration data for radius-flow-tap policy rule mapping.

Options ***address***— IPv4 or IPv6 address for the radius-flow-tap policy.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

 Copyright © 2013, Juniper Networks, Inc.

source-count (Dynamic MLD Interface)

Syntax source-count *number*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the number of multicast source addresses that should be accepted for each static group created on dynamic interfaces.

Options *number*—Number of source addresses.

Default: 1

Range: 1 through 1024

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Enabling MLD Static Group Membership*

source-gateway (Tunnel Profile)

Syntax

```
source-gateway {
   address client-ip-address;
   gateway-name client-name;
}
```

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IP address and hostname of the source gateway at the local L2TP tunnel endpoint, the LAC.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 403

source-increment (Dynamic MLD Interface)

Syntax source-increment *increment*;

Hierarchy Level [edit dynamic-profile *profile-name* protocols mld interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the number of times the address should be incremented for each static group created on the dynamic interface. The increment is specified in a format similar to an IPv6 address.

Options ***increment***—Number of times the source address should be incremented.

Default: ::1

Range: ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Enabling MLD Static Group Membership*

source-ipv4-address

Syntax source-ipv4-address *ipv4-address*;

Hierarchy Level [edit services radius-flow-tap]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the source IP address used in the IP header that is prepended to mirrored packets sent to a mediation device.

Options ***ipv4-address***—IPv4 address.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

 Copyright © 2013, Juniper Networks, Inc.

source-port (Subscriber Secure Policy)

Syntax source-port *port-number*;

Hierarchy Level [edit services radius-flow-tap policy *policy-name* inet drop-policy *rule-name* from],
[edit services radius-flow-tap policy *policy-name* inet6 drop-policy *rule-name* from]

Release Information Statement introduced in Junos OS Release 12.3.

Description Specify the match source port for the radius-flow-tap policy.

Options ***port-number***— Number of the IPv4 or IPv6 source port for the radius-flow-tap policy.

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 on page 1245
- Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview on page 1261

spi

Syntax

```
spi hexadecimal-value {
   algorithm (hmac-md5 | md5);
   entity-type (host | mobility-agent);
   key (hex | ascii) string;
   replay-method (none | timestamp seconds);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip peer ip-address *address*],
[edit logical-systems *logical-system-name* services mobile-ip peer nai *user@domain*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer ip-address *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain*],
[edit routing-instances *routing-instances-name* services mobile-ip peer ip-address *address*],
[edit routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain*],
[edit services mobile-ip peer ip-address *address*],
[edit services mobile-ip peer nai *user@domain*]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* services mobile-ip peer ip-address *address*]**, **[edit logical-systems *logical-system-name* services mobile-ip peer nai *user@domain*]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer ip-address *address*]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain*]**, **[edit routing-instances *routing-instances-name* services mobile-ip peer ip-address *address*]**, and **[edit routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain*]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the security parameter index for identifying a security context between a pair of nodes among the contexts available in the Mobility Security Association. The index selects the authentication algorithm and key.

Options

hexadecimal-value—Security parameter index identifier.

Range: 100 to FFFFFFFF

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

Copyright © 2013, Juniper Networks, Inc.

ssm-map (Dynamic IGMP Interface)

Syntax ssm-map *ssm-map-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Apply an SSM map to an IGMP interface.

Options ***ssm-map-name***—Name of SSM map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about configuring SSM maps, see "Source-Specific Multicast Groups Overview" in the *Multicast Protocols Feature Guide for Routing Devices*

ssm-map (Dynamic MLD Interface)

Syntax ssm-map *ssm-map-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Apply an SSM map to a dynamic MLD interface.

Options ***ssm-map-name***—Name of SSM map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring SSM Mapping*

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 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65

 Copyright © 2013, 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.

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 323

static (Dynamic IGMP Interface)

Syntax

```
static {
  group group;
  group group {
    source source;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Test multicast forwarding on an interface without a receiver host.

Options

The remaining statements are explained separately.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about testing multicast forwarding without a receiver host, see "Enabling IGMP Static Group Membership" in the *Multicast Protocols Feature Guide for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

static (Dynamic MLD Interface)

Syntax

```
static {
  group multicast-group-address {
    exclude;
    group-count number;
    group-increment increment;
    source ip-address {
      source-count number;
      source-increment increment;
    }
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Test multicast forwarding on an interface.

The remaining statements are explained separately.

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- *Enabling MLD Static Group Membership*

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.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring Subscribers over Static Interfaces on page 500

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.
Option **volume-time** introduced 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. This option is not available for Mobile IP.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Mobile IP Home Agent Elements and Behavior on page 557
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24

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 Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- Preventing Binding of Clients That Do Not Support Reconfigure Messages on page 249

Copyright © 2013, 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 189
- Configuring Domain Name Usage for Domain Maps on page 187

subscriber-identification (PTSP)

Syntax subscriber-identification *subscriber-identification*

Hierarchy Level [edit system services packet-triggered-subscribers partition radius *radius-partition-name*]

Release Information Statement introduced in Junos OS Release 11.3.

Description Configure the subscriber identification to be used in a PTSP partition. You can configure the subscriber identification only in a RADIUS partition.

Options ***subscriber-identification***—String of user-defined characters or a RADIUS attribute type that is supported by the PTSP application. To enable subscriber identification for the specified RADIUS attribute, you may configure the following RADIUS attributes:

- $attribute-1$—User-Name
- $attribute-4$—NAS-IP-Address
- $attribute-5$—NAS-Port
- $attribute-8$—Framed-IP-Address
- $attribute-32$—NAS-Identifier
- $attribute-87$—NAS-Port-ID

When configuring subscriber identification, you must precede the "$" with a slash (\) to enable the CLI interface to process and store the variable correctly.

NOTE: **The IP address is formatted in dotted decimal notation—for example, 192.168.1.1. All the other numeric values are converted to a string of characters.**

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the PTSP Partition on page 526

 Copyright © 2013, Juniper Networks, Inc.

subscriber-management (Subscriber Management)

Syntax

```
subscriber-management {
  enforce-strict-scale-limit-license;
  gres-route-flush-delay;
  maintain-subscriber {
    interface-delete;
  }
  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 and tracing operations.

The remaining statements are explained separately.

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 on page 235
- Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 693
- Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events on page 232
- Tracing Subscriber Management Database Operations for Subscriber Access on page 158

Copyright © 2013, Juniper Networks, Inc.

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.

swap (Dynamic VLANs)

Syntax swap;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 10.4.

Description For dynamic VLAN interfaces, specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Rewriting the VLAN Tag on Tagged Frames*
- *Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution*

 Copyright © 2013, Juniper Networks, Inc.

tag (Access)

Syntax tag *route-tag*;

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*],
[edit dynamic-profiles *profile-name* routing-options access route *prefix*]

Release Information Statement introduced in Junos OS Release 10.2.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*]** hierarchy levels introduced in Junos OS Release 10.1.

Description Dynamically configure the tag for an access route.

Options ***route-tag***—Either the specific tag you want to assign to the access route or either of the following tag variables:

- **$junos-framed-route-tag**—Tag assigned to an IPv4 access route; the variable is dynamically replaced with the preference value (Subattribute 6) from the RADIUS Framed-Route attribute [22].
- **$junos-framed-route-ipv6-tag**—Tag assigned to an IPv6 access route; the variable is dynamically replaced with the preference value (Subattribute 6) from the RADIUS Framed-IPv6-Route attribute [99].

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Access Routes for Subscriber Management on page 690

tag (Dynamic Profiles)

Syntax tag *tag-number*;

Hierarchy Level [edit dynamic-profiles *profile-name* variables radius vendor-id]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure a tag for a RADIUS attribute as a variable in a dynamic profile.

Options ***tag-number***—Tag number for the RADIUS attribute.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring User-Defined CoS Variables in a Dynamic Service Profile on page 1000

tag-protocol-id (Dynamic VLANs)

Syntax tag-protocol-id *tpids*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 10.4.

Description For dynamic VLAN interfaces, configure the outer TPID value. All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level.

Default If the **tag-protocol-id** statement is not configured, the TPID value is 0x8100.

Options ***tpids***—TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal format.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring Inner and Outer TPIDs and VLAN IDs*

 Copyright © 2013, 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. The target-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.

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 186

Copyright © 2013, Juniper Networks, Inc.

target-routing-instance (Domain Map)

Syntax target-routing-instance *routing-instance-name*;

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.

Description Configure a non-default routing instance for the target logical system for the subscriber's interface in a domain map.

- When configured at the **[edit access domain map *domain-map-name*]** hierarchy level, this statement configures the routing instance used with the default target logical system.
- When configured at the **[edit access domain map *domain-map-name* target-logical-system *logical-system-name*]** hierarchy level, this statement configures the routing instance used with the specified non-default target logical system.

Default Default routing instance for the subscriber.

Options ***routing-instance-name***—Name of the routing instance.

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 186

 Copyright © 2013, Juniper Networks, Inc.

targeted-distribution (Static Interfaces over Aggregated Ethernet)

Syntax targeted-distribution;

Hierarchy Level [edit interfaces demux0 unit *logical-unit-number*],
[edit interfaces pp0 unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure egress data for a logical interface to be sent across a single member link in an aggregated Ethernet bundle. A backup link is provisioned with CoS scheduling resources in the event that the primary assigned link goes down. The aggregated Ethernet interface must be configured without link protection.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- CoS for PPPoE Subscriber Interfaces Overview on page 1014
- Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces on page 833
- Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface on page 833
- *Targeted Traffic Distribution on Aggregated Ethernet Interfaces in a Virtual Chassis*
- *Configuring Module Redundancy for a Virtual Chassis*
- *Configuring Chassis Redundancy for a Virtual Chassis*

Copyright © 2013, Juniper Networks, Inc.

term

Syntax

```
term term-name {
  from {
    match-conditions;
  }
  then {
    action;
    action-modifiers;
  }
  only-at-create;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall family *family-name* fast-update-filter *filter-name*],
[edit dynamic-profiles *profile-name* firewall family *family-name* filter *filter-name*]

Release Information

Statement introduced in Junos OS Release 9.6.
Support at the **[edit dynamic-profiles ... filter *filter-name*]** hierarchy level introduced in Junos OS Release 11.4.

Description

Define terms for fast update filters.

Options

action—(Optional) An action to take if conditions match. If you do not specify an action, the packets that match the conditions in the **from** statement are accepted.

action-modifiers—(Optional) One or more actions to perform on a packet.

from—(Optional) Match packet fields to values. If not included, all packets are considered to match and the actions and action modifiers in the **then** statement are taken.

match-conditions—One or more conditions to make a match.

only-at-create—(Optional) Specify that the term is added only when the fast update filter is first created. No subsequent changes can be made to the term in the filter. Use this option only for terms that do not include subscriber-specific data in their match conditions, such as common or default terms (for example, counting the default drop packets).

term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

then—(Optional) Actions to take on matching packets. If not included and a packet matches all the conditions in the **from** statement, the packet is accepted.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Fast Update Filters on page 1183
- Configuring Terms for Fast Update Filters on page 1185

Copyright © 2013, Juniper Networks, Inc.

- Fast Update Filter Match Conditions on page 1186
- Fast Update Filter Actions and Action Modifiers on page 1187

term (Captive Portal Content Delivery)

Syntax

```
term term-name {
  from {
    application [ application-name];
    destination-address address <except>;
    destination-prefix-list list-name <except>;
  }
  then {
    action;
    action-modifiers;
  }
}
```

Hierarchy Level

[edit services (captive-portal-content-delivery) captive-portal-content-delivery rule *rule-name*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Define the captive-portal-content-delivery term properties.

Options

term-name—Identifier for the term.

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

- Redirecting HTTP Requests Overview on page 1225

terminate-code

Syntax

```
terminate-code {
   (aaa (deny | shutdown) | dhcp | l2tp | ppp) term-reason radius term-cause;
}
```

Hierarchy Level

[edit access]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Customize mappings between a terminate reason and a RADIUS Acct-Terminate-Cause attribute.

Options

aaa—Defines AAA as the protocol type, and specifies either the **deny** or **shutdown** action to perform with the specified terminate reason.

dhcp—Defines DHCP as the protocol type associated with the specified terminate reason.

l2tp—Defines L2TP as the protocol type associated with the specified terminate reason.

ppp—Defines PPP as the protocol type associated with the specified terminate reason.

term-reason—Terminate reason for the specified protocol type.

term-cause—Standards-based RADIUS Acct-Terminate-Cause code that identifies the terminate reason.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Custom Terminate Reason Mappings on page 148
- AAA Terminate Reasons on page 123
- DHCP Terminate Reasons on page 124
- L2TP Terminate Reasons on page 124
- PPP Terminate Reasons on page 140

Copyright © 2013, Juniper Networks, Inc.

tftp-server

Syntax tftp-server *ip-address*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes]

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 168

Copyright © 2013, Juniper Networks, Inc.

then

Syntax

```
then {
  action;
  <action-modifiers>;
}
```

Hierarchy Level

[edit services (captive-portal-content-delivery) captive-portal-content-delivery rule *rule-name* term *term-name*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Define the captive-portal-content-delivery term actions and any optional action modifiers.

Options

action—Actions to accept, redirect, or rewrite packets and all subsequent packets in flows that match the rules.

- **accept**—Accept the packets and all subsequent packets in flows that match the rules.
- **redirect**—Redirect the packet and all subsequent packets in flows that match the rules. You can optionally configure the *url* action modifier.
- **rewrite**— Rewrite the packet and all subsequent packets in flows that match the rules. You can optionally configure the *destination-address* and *destination-port* action modifiers.
- **syslog**— System log information about the packet.

***action-modifiers* (Optional)**—Additional actions to accept, redirect, or rewrite packets and all subsequent packets in flows that match the rules.

- *destination-address* —(Optional) Destination address of the rewrite packet.
- *destination-port* —(Optional) Destination address and destination port of the rewrite packet.
- *url*—(Optional) URL of the redirect packet.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Redirecting HTTP Requests Overview on page 1225
- *Routing Policy Feature Guide for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

three-color-policer (Configuring)

Syntax

```
three-color-policer policer-name {
  action {
    loss-priority high then discard;
  }
  filter-specific;
  logical-interface-policer;
  physical-interface-policer;
  shared-bandwidth-policer;
  single-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    peak-burst-size bytes;
    peak-information-rate bps;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall],
[edit firewall],
[edit logical-systems *logical-system-name* firewall]

Release Information

Statement introduced before Junos OS Release 7.4.
The **action** and **single-rate** statements added in Junos OS Release 8.2.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... firewall]** hierarchy level introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure a three-color policer.

Options

policer-name—Name of the three-color policer. Reference this name when you apply the policer to an interface.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Statement Hierarchy for Configuring Policers*
- *Configuring Tricolor Marking Policers*
- *Three-Color Policer Configuration Guidelines*
- *Basic Single-Rate Three-Color Policers*

Copyright © 2013, Juniper Networks, Inc.

- *Basic Two-Rate Three-Color Policers*
- *Two-Color and Three-Color Logical Interface Policers*
- *Two-Color and Three-Color Physical Interface Policers*
- *Two-Color and Three-Color Policers at Layer 2*

threshold (detection-time)

Syntax threshold *milliseconds*;

Hierarchy Level
[edit system services dhcp-local-server liveness-detection method bfd detection-time],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd detection-time],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd detection-time],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd detection-time]

Release Information
Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

NOTE: **The threshold time must be greater than or equal to the minimum-interval or the minimum-receive-interval.**

Options ***milliseconds***— Value for the detection time adaptation threshold.

Range: 1 through 255,000

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

 Copyright © 2013, Juniper Networks, Inc.

threshold (transmit-interval)

Syntax threshold *milliseconds*;

Hierarchy Level
[edit system services dhcp-local-server liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd transmit-interval]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify the threshold for detecting the adaptation of the transmit interval. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.

Options ***milliseconds*** — Threshold value.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

NOTE: **The threshold value specified in the threshold statement must be greater than the value specified in the minimum-interval statement for the transmit-interval statement.**

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

Copyright © 2013, 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 Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 246

Copyright © 2013, 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 a RADIUS server.

Options ***seconds***—Amount of time to wait.

Range: 1 through 90 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 23
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24
- *Example: Configuring CHAP Authentication with RADIUS*
- *Configuring RADIUS Authentication for L2TP*

Copyright © 2013, Juniper Networks, Inc.

timestamp-tolerance

Syntax timestamp-tolerance *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit services mobile-ip home-agent virtual-network home-agent-address *ip-address*]

Release Information
Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Configure the acceptable difference between a registration request timestamp and the local time of the home agent.

Options
timestamp-tolerance *seconds*—Acceptable difference in time between a registration request timestamp and the local time of the home agent.

Range: 1 through 255 seconds

Default: 7 seconds

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

Copyright © 2013, 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.

Related Documentation

- Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- Configuring a Token for DHCP Local Server Authentication on page 248

tos-reflect (L2TP)

Syntax tos-reflect;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the LNS to reflect the IP ToS value in the inner IP header to the outer IP header. When CoS rewrite rules are also configured (**[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]**), the rewrite is performed only on the inner IP ToS; the outer IP TOS value is not affected.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Access Profiles for L2TP Tunnel Groups*
- Configuring Dynamic CoS for an L2TP LNS Inline Service on page 1026

 Copyright © 2013, Juniper Networks, Inc.

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 255
- Tracing Extended DHCP Operations for Specific Interfaces on page 260

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.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing Extended DHCP Operations on page 255
- Tracing Extended DHCP Operations for Specific Interfaces on page 260

Copyright © 2013, 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>;
  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

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **address-assignment**—Trace address-assignment pool events
- **all**—Trace all tracing operations
- **configuration**—Trace configuration events
- **framework**—Trace authentication framework events
- **gx-plus**—Trace Gx-Plus events
- **jsrc**—Trace JSRC events
- **ldap**—Trace LDAP authentication events
- **local-authentication**—Trace local authentication events
- **radius**—Trace RADIUS authentication events
- **user-access**—Trace user access events, such as login, logout, and authenticate.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Tracing General Authentication Service Processes on page 175

Copyright © 2013, 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 processes.

Options

file *filename*— Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. Include the **disable** option after a flag to disable tracing for that flag. You can include the following flags:

- **all**—Trace all operations.
- **config**—Trace configuration events.
- **cos**—Trace class-of-service events.
- **general**—Trace general flow.
- **packet**—Trace ANCP packet transmit and receive operations.
- **process**—Trace process internals.
- **protocol**—Trace protocol events.
- **restart**—Trace process restart flow
- **routing-socket**—Trace routing socket events.
- **session**—Trace connection events and flow.
- **startup**—Trace ANCP startup events and flow.
- **subscriber**—Trace subscriber events.
- **timer**—Trace timer processing.

level—Level of tracing to perform. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions reqcliring special handling.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

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 Operations for Subscriber Access on page 1345

Copyright © 2013, Juniper Networks, Inc.

traceoptions (Captive Portal Content Delivery)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag configuration;
  flag general;
  flag gres;
  flag rtsock;
  flag statistics;
  flag "all";
  no-remote-trace;
}
```

Hierarchy Level

[edit services (captive-portal-content-delivery)captive-portal-content-delivery]

Release Information

Statement introduced in Junos OS Release 10.4.
Support at the **[edit services captive-portal-content-delivery]** hierarchy level introduced in Junos OS Release 10.4.

Description

Define tracing operations for captive-portal-content-delivery 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**. Ensure that filenames are unique for each logical system or routing instance in which Mobile IP is configured.

NOTE: Global messages (common to all logical systems and routing instances) are always saved in /var/log/mipd. Messages that are specific to a logical system or routing instance are never saved in /var/log/mipd. If you do not configure a trace filename for a logical system or routing instance, then nothing is traced for that entity.

size ***size***—(Optional) Maximum size of each trace file, in 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: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

files ***number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. 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.
- **configuration**—Trace home agent state machine operations.
- **general**—Trace general operations.
- **gres**—Trace graceful routing switchover operations.
- **rtsock**—Trace routing socket operations.
- **statistics**—Trace statistics operations.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Redirecting HTTP Requests Overview on page 1225

Copyright © 2013, 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.

Replaces 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.
- **packet**—Trace packet and option decoding operations.
- **performance**—Trace performance measurement operations.
- **profile**—Trace profile operations.
- **rpd**—Trace routing protocol process events.

- **rtsock**—Trace routing socket operations.
- **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.

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 (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 Extended DHCP Operations on page 255

Copyright © 2013, 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.
- **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 475

Copyright © 2013, Juniper Networks, Inc.

traceoptions (L2TP)

Syntax

```
traceoptions {
  debug-level level;
  file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  filter {
    protocol name;
    user-name username;
  }
  flag flag;
  interfaces interface-name {
    debug-level level;
    flag flag;
  }
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

Hierarchy Level [edit services l2tp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define tracing operations for L2TP processes.

Options **debug-level** ***level***—Trace level for PPP, L2TP, RADIUS, and UDP; this option does not apply to L2TP on MX Series routers:

- **detail**—Trace detailed debug information.
- **error**—Trace error information.
- **packet-dump**—Trace packet decoding information.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

filter protocol *name*—Additional filter for the specified protocol; this option does not apply to L2TP on MX Series routers:

- **l2tp**
- **ppp**
- **radius**
- **udp**

filter user-name *username*—Additional filter for the specified username; this option does not apply to L2TP on MX Series routers.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **configuration**—Trace configuration events.
- **events**—Trace interface events.
- **general**—Trace general events.
- **gres**—Trace GRES events.
- **init**—Trace daemon initialization.
- **ipc-rx**—Trace IPC receive events.
- **ipc-tx**—Trace IPC transmit events.
- **memory**—Trace memory management code.
- **message**—Trace message processing code.
- **packet-error**—Trace packet error events.
- **parse**—Trace parsing events.
- **protocol**—Trace L2TP events.
- **receive-packets**—Trace received L2TP packets.
- **routing-process**—Trace routing process interactions.
- **routing-socket**—Trace routing socket events.
- **session-db**—Trace session database interactions.
- **states**—Trace state machine events.
- **timer**—Trace timer events.
- **transmit-packets**—Trace transmitted L2TP packets.
- **tunnel**—Trace tunnel events.

 Copyright © 2013, Juniper Networks, Inc.

interfaces ***interface-name***—Apply L2TP traceoptions to a specific services interface. This option does not apply to L2TP on MX Series routers.

- **debug-level** ***level***—Trace level for the interface; this option does not apply to L2TP on MX Series routers:
  - **detail**—Trace detailed debug information.
  - **error**—Trace error information.
  - **extensive**—Trace all PIC debug information.
- **flag** ***flag***—Tracing operation to perform for the interface. This option does not apply to L2TP on MX Series routers. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:
  - **all**—Trace everything.
  - **ipc**—Trace L2TP Inter-Process Communication (IPC) messages between the PIC and the Routing Engine.
  - **packet-dump**—Dump each packet content based on debug level.
  - **protocol**—Trace L2TP, PPP, and multilink handling.
  - **system**—Trace packet processing on the PIC.

level—Specify level of tracing to perform. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- **all**—Match messages of all levels.
- **error**—Match error messages.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- *Tracing L2TP Operations*
- Tracing L2TP Operations for Subscriber Access on page 434

Copyright © 2013, Juniper Networks, Inc.

traceoptions (Mobile IP)

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* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip],
[edit routing-instances *routing-instances-name* services mobile-ip],
[edit services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define tracing operations for Mobile IP 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**. Ensure that filenames are unique for each logical system or routing instance in which Mobile IP is configured.

NOTE: Global messages (common to all logical systems and routing instances) are always saved in **/var/log/mipd**. Messages that are specific to a logical system or routing instance are never saved in **/var/log/mipd**. If you do not configure a trace filename for a logical system or routing instance, then nothing is traced for that entity.

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 operations.
- **binding**—Trace bindings.

- **event**—Trace events.
- **ha-fsm**—Trace home agent state machine operations.
- **home-agent**—Trace home agent operations.
- **interface-database**—Trace interface database operations.
- **packet**—Trace packet decoding operations.
- **protocol**—Trace protocol operations.
- **rtsock**—Trace routing socket operations.
- **session-db**—Trace session database events.
- **signal**—Trace signal operations.
- **subscriber**—Trace subscriber events.
- **timer**—Trace timer events.
- **trace**—Trace changes in tracing.
- **tunnel**—Trace tunneling operations.
- **user-interface**—Trace user interface events.

level—Specify level of tracing to perform. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Copyright © 2013, 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 Mobile IP Operations for Subscriber Access on page 573

traceoptions (PTSP)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag <disable>;
  no-remote-trace;
}
```

Hierarchy Level [edit system services packet-triggered-subscribers]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define tracing operations for PTSP.

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.
- **configuration**—Trace configuration events.
- **general**—Trace general flow.
- **peer**—Trace SRC peer events.
- **pic**—Trace PIC events.
- **rtsock**—Trace routing socket events.
- **session**—Trace session events.

disable—Disable this trace flag.

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.

 Copyright © 2013, Juniper Networks, Inc.

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 Packet-Triggered Subscriber Operations on page 527

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:

- **all**—Match all levels.

Copyright © 2013, 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—(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 509

traceoptions (Subscriber Management)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
}
```

Hierarchy Level

[edit system services subscriber-management]

Release Information

Statement introduced in Junos OS Release 11.1.

Description

Define tracing operations for subscriber management interface processes.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **database**—Trace database events.
- **general**—Trace general events.
- **issu**—Trace unified ISSU events.
- **server**—Trace server events.
- **session-db**—Trace session database interactions.
- **ui**—Trace user interface events.

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

Default: 128 KB

Copyright © 2013, Juniper Networks, Inc.

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 158

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.

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

 Copyright © 2013, Juniper Networks, Inc.

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 161

traffic-control-profiles (Dynamic CoS Definition)

Syntax

```
traffic-control-profiles profile-name {
  adjust-minimum rate;
  delay-buffer-rate (percent percentage | rate);
  excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
  excess-rate-high (percent percentage | proportion value);
  excess-rate-low (percent percentage | proportion value);
  guaranteed-rate (percent percentage | rate) <burst-size bytes>;
  overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
  scheduler-map map-name;
  shaping-rate (percent percentage | rate | predefined-variable) <burst-size bytes>;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Configure traffic shaping and scheduling profiles.

Options

profile-name—Name of the traffic-control profile.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Traffic Scheduling and Shaping for Subscriber Access on page 971
- Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers on page 1125
- **output-traffic-control-profile on page 1839**

transmit-interval

Syntax

```
transmit-interval {
  threshold milliseconds;
  minimum-interval milliseconds;
}
```

Hierarchy Level

[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the Bidirectional Forwarding Detection (BFD) transmit interval.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

Copyright © 2013, Juniper Networks, Inc.

transmit-rate (Dynamic Schedulers)

Syntax

transmit-rate (*rate* | percent *percentage* | remainder | percent *percentage*
$junos-cos-scheduler-tx) <exact | rate-limit>;

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]

Release Information

Statement introduced in Junos OS Release 9.3.
The **$junos-cos-scheduler-tx** predefined variable introduced in Junos OS Release 9.4.

Description

Specify the transmit rate or percentage for a scheduler in a dynamic profile.

Default

If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

Options

rate—Transmission rate, in bps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 3200 through 160,000,000,000 bps

percent *percentage*—Percentage of transmission capacity. A percentage of zero drops all packets in the queue.

Range: 0 through 100 percent

remainder—Use remaining rate available.

$junos-cos-scheduler-tx—Junos predefined variable that is replaced with the transmission rate obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

exact—(Optional) Enforce the exact transmission rate. Under sustained congestion, a rate-controlled queue that goes into negative credit fills up and eventually drops packets. Make sure this value never exceeds the rate-controlled amount.

rate-limit—(Optional) Limit the transmission rate to the rate-controlled amount during congestion. In contrast to the **exact** option, when there is no congestion, the scheduler with the **rate-limit** option shares unused bandwidth above the rate-controlled amount.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Configuring Schedulers in a Dynamic Profile for Subscriber Access on page 973
- **scheduler on page 1968**

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.

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 471
- Configuring the Diameter Transport on page 474

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 471
- Configuring Diameter Peers on page 473

 Copyright © 2013, 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.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 245
- Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 247
- radius-disconnect on page 1919

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 296
- Overriding the Default DHCP Relay Configuration Settings on page 293

Copyright © 2013, Juniper Networks, Inc.

tunnel (L2TP)

Syntax

```
tunnel {
  assignment-id-format (assignment-id | client-server-id);
  idle-timeout seconds;
  retransmission-count-established count;
  retransmission-count-not-established count;
  rx-window-size
  tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port);
}
```

Hierarchy Level

[edit services l2tp]

Release Information

Statement introduced in Junos OS Release 11.4.
rx-window-size option introduced in Junos OS Release 13.2.

Description

Configure L2TP tunnel characteristics.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an L2TP LAC on page 402
- Configuring an L2TP LNS with Inline Service Interfaces on page 417

tunnel (Tunnel Profile)

Syntax

```
tunnel tunnel-id {
  identification name;
  logical-system logical-system-name;
  max-sessions number;
  medium type;
  nas-port-method cisco-avp;
  preference number;
  remote-gateway {
    address server-ip-address;
    gateway-name server-name;
  }
  routing-instance routing-instance-name;
  secret password;
  source-gateway {
    address client-ip-address;
    gateway-name client-name;
  }
  type tunnel-type;
}
```

Hierarchy Level

[edit access tunnel-profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Define the attributes of a tunnel for the tunnel profile. You can define up to 31 tunnels for each tunnel profile.

Options

tunnel-id—Unique integer that identifies a tunnel defined within a profile. For a subscriber, RADIUS attributes and VSAs can supply or override the attributes defined here for the tunnel.

Range: 1 through 31

The remaining statements are explained separately.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

Copyright © 2013, Juniper Networks, Inc.

tunnel-group

Syntax

```
tunnel-group group-name {
  aaa-access-profile profile-name;
  dynamic-profile profile-name;
  hello-interval seconds;
  hide-avps;
  l2tp-access-profile profile-name;
  local-gateway address address;
  maximum-send-window packets;
  ppp-access-profile profile-name;
  receive-window packets;
  retransmit-interval seconds;
  service-device-pool pool-name;
  service-interface interface-name;
  syslog {
    host hostname {
      services severity-level;
      facility-override facility-name;
      log-prefix prefix-value;
    }
  }
  tos-reflect;
  tunnel-switch-profile profile-name;
  tunnel-timeout seconds;
}
```

Hierarchy Level

[edit services l2tp]

Release Information

Statement introduced before Junos OS Release 7.4.
Support for MX Series routers introduced in Junos OS Release 11.4.

Description

Specify the L2TP tunnel properties.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

group-name—Identifier for the tunnel group.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring L2TP Tunnel Groups*
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 429

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 189
- Configuring a Tunnel Profile for Subscriber Access on page 403

tunnel-profile (L2TP Tunnel Switching)

Syntax tunnel-profile *profile-name*;

Hierarchy Level [edit access tunnel-switch-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the name of an L2TP tunnel profile that defines the tunnel to which PPP subscriber traffic is switched.

Options *profile-name*—Unique name that identifies the tunnel profile; configured with the **tunnel-profile** statement at the **[edit access]** hierarchy level.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring L2TP Tunnel Switching on page 432

 Copyright © 2013, Juniper Networks, Inc.

tunnel-profile (Tunnel Profile)

Syntax

```
tunnel-profile profile-name {
  tunnel tunnel-id {
    identification name;
    logical-system logical-system-name;
    max-sessions number;
    medium type;
    nas-port-method cisco-avp;
    preference number;
    remote-gateway {
      address server-ip-address;
      gateway-name server-name;
    }
    routing-instance routing-instance-name;
    secret password;
    source-gateway {
      address client-ip-address;
      gateway-name client-name;
    }
    type tunnel-type;
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 10.4.

Description Define a tunnel profile for subscriber access.

Options *profile-name*—Unique name that identifies the tunnel profile. The profile can be referenced from within a domain map or by the RADIUS Tunnel-Group VSA [26-64].

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

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

Syntax tunnel-switch-profile *profile-name*;

Hierarchy Level [edit access domain map *domain-map-name*],
[edit services l2tp],
[edit services l2tp tunnel-group *group-name*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify a tunnel switch profile that determines whether packets in an L2TP session from a LAC are switched to another session that has a different destination LNS.

Options *profile-name*—Unique name that identifies the tunnel switch profile.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring L2TP Tunnel Switching on page 432

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

Syntax

```
tunnel-switch-profile profile-name {
  avp {
    bearer-type action;
    calling-number action;
    cisco-nas-port-info action;
  }
  tunnel-profile profile-name;
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 13.2.

Description Define a tunnel switch profile for subscriber access; specify actions to take for L2TP AVPs in the switched packets and the profile for the tunnel to which the PPP traffic is switched.

The remaining statements are explained separately.

Options *profile-name*—Unique name that identifies the tunnel switch profile. The profile can be applied as a default or referenced from within a domain map, a tunnel group, or by the RADIUS Tunnel-Group VSA [26-64].

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring L2TP Tunnel Switching on page 432

 Copyright © 2013, Juniper Networks, Inc.

two-rate

Syntax

```
two-rate {
  (color-aware | color-blind);
  committed-information-rate bps;
  committed-burst-size bytes;
  peak-information-rate bps;
  peak-burst-size bytes;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall three-color-policer *name*],
[edit firewall three-color-policer *policer-name*],
[edit logical-systems *logical-system-name* firewall three-color-policer *policer-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... three-color-policer *name*** hierarchy levels introduced in Junos OS Release 11.4.

Description

Configure a two-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), peak information rate (PIR), and peak burst size (PBS).

Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the PIR or the PBS are assigned medium-high loss priority (yellow). Packets that exceed the PIR and the PBS are assigned high loss priority (red).

Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Three-Color Policer Configuration Overview*
- color-aware on page 1514
- color-blind on page 1515
- single-rate on page 1991

Copyright © 2013, Juniper Networks, Inc.

tx-address-change (L2TP LAC)

Syntax tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port);

Hierarchy Level [edit services l2tp tunnel]

Description Configure whether the LAC accepts or ignores requests from a peer to change the destination IP address, UDP port, or both.

Default The LAC accepts the IP address change from its peer.

Options

accept—Accept a change in IP address , UDP port, or both.

ignore—Ignore a change request for IP address or UDP port.

ignore-ip-address—Ignore a change request for IP address.

ignore-udp-port—Ignore a change request for UDP port.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the LAC to Ignore Address and Port Changes Requested by the LNS on page 413

 Copyright © 2013, Juniper Networks, Inc.

tx-connect-speed-method (L2TP LAC)

Syntax tx-connect-speed-method *method*;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 11.4.
Options **ancp**, **pppoe-ia-tag**, and **static** introduced in Junos OS Release 13.1.

Description Specify the method that determines the connection speed values sent from the LAC to the LNS in Incoming-Call-Connected (ICCN) messages. The transmit speed is conveyed in AVP 24 (Tx Connect Speed) and the receive speed is conveyed in AVP 38 (Rx Connect Speed). Both values are in bits per seconds (BPS).

Default static

Options *method*—Method used to derive the connection speed values.

- **ancp**—The speed is derived from the configured ANCP value for the underlying interface. You can change this rate after a subscriber has logged in, but those changes do not affect the actual rate used by the LNS. The **ancp** method gets the highest preference among the methods configured.
- **pppoe-ia-tags**—PPPoE IA tags sent from the DSLAM to the LAC. This speed value transmitted when a subscriber logs in and it cannot be subsequently changed. This value is used when the **ancp** value is not available. This speed does not apply to the subscribers that are already logged in; it is effective only for subscribers that log in after this setting has been saved. AVP 24 is the value of Actual-Data-Rate-Downstream (VSA 26-129). AVP 38 is the value of Actual-Data-Rate-Upstream (VSA 26-130), and is sent only when the VSA values differ.

  NOTE: **This rate does not affect the subscribers already logged in; however, new subscribers inherit the new rate.**

- **static**—The speed is derived from the recommended (advisory) shaping rate configured on the PPPoE logical interface underlying the subscriber interface. If the advisory speed is not configured on the underlying interface, then the actual port speed is used. The default method, when no other methods yield a value, is the **static** method or the advisory speed method. If the advisory speed is not configured, then the actual port speed is used. For ge and xe interfaces, the speed value is set to 10,000,000 and for ae interfaces, the speed value is set to 0 and sent in both AVP 24 and AVP 38.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method to Set the LAC Connection Speeds to the LNS on page 408

type (Tunnel Profile)

Syntax

type *tunnel-type*;

Hierarchy Level

[edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Specify the tunnel type (protocol).

Default

l2tp

Options

tunnel-type—Tunnel protocol type. The only value currently available is **l2tp**.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 403

Copyright © 2013, Juniper Networks, Inc.

t1-percentage

Syntax t1-percentage *percentage*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify a percentage of the preferred-lifetime value. After this percentage of time elapses, the router requests an extension on its lease from the originating DHCP server.

The t1-percentage value must be less than the t2-percentage value or a commit error occurs.

Options *percentage*—Percentage of the preferred-lifetime value.

Range: 0 through 100

Default: Default depends on the value of preferred-lifetime.

- When the preferred-lifetime value is finite and the t1-percentage value is not configured, the default is 50 percent.
- When the preferred-lifetime value is infinite, the default is 100 percent.

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 168
- t2-percentage on page 2072
- preferred-lifetime on page 1889
- dhcp-attributes (Address-Assignment Pools) on page 1542
- DHCP Attributes for Address-Assignment Pools on page 173

Copyright © 2013, Juniper Networks, Inc.

t2-percentage

Syntax t2-percentage *percentage*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify a percentage of the preferred-lifetime value. After this percentage of time elapses, the router requests an extension on its lease from any available DHCP server.

The t2-percentage value must be greater than the t1-percentage value or a commit error occurs.

Options *percentage*—Percentage of the preferred-lifetime value, at which time you want the client (router) to contact the DHCPv6 server to extend the IPv6 configuration.

Range: 0 through 100

Default: Default depends on the value of the preferred-lifetime statement.

- When the preferred-lifetime value is finite and the t2-percentage value is not configured, the default is 80 percent.
- When the preferred-lifetime value is infinite, the default is 100 percent.

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 168
- t1-percentage on page 2071
- preferred-lifetime on page 1889
- DHCP Attributes for Address-Assignment Pools on page 173
- dhcp-attributes (Address-Assignment Pools) on page 1542

 Copyright © 2013, Juniper Networks, Inc.

uid

Syntax uid;

Hierarchy Level [edit dynamic-profiles *profile-name* variables *variable-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure a unique ID for parameterized filters in a dynamic profile created for services. The values that the system uses for these variables are applied when the subscriber authenticates.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Dynamic Variables Overview on page 640

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 ANCP on page 1332
- Associating an Access Node with Subscribers for ANCP Operations on page 1334

underlying-interface (demux0)

Syntax underlying-interface *underlying-interface-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces demux0 *interface-name* unit unit *logical-unit-number* demux-options]

Release Information Statement introduced in Junos OS Release 9.3.
Support for aggregated Ethernet introduced in Junos OS Release 9.4.

Description Configure the underlying interface on which the demultiplexing (demux) interface is running.

Options ***underlying-interface-name***—Either the specific name of the interface on which the DHCP discover packet arrives or one of the following interface variables:

- **$junos-underlying-interface** when configuring dynamic IP demux interfaces.
- **$junos-interface-ifd-name** when configuring dynamic VLAN demux interfaces.

The variable is used to specify the underlying interface when a new demux interface is dynamically created. The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.

NOTE: **Logical demux interfaces are currently supported on Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interfaces.**

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static Subscriber Interfaces Using IP Demux Interfaces on page 766
- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- Configuring Static Subscriber Interfaces Using VLAN Demux Interfaces on page 767
- Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles on page 772
- For information about static underlying interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

 Copyright © 2013, Juniper Networks, Inc.

underlying-interface (Dynamic PPPoE)

Syntax underlying-interface *interface-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" pppoe-options]

Release Information Statement introduced in Junos OS Release 10.1.

Description In a dynamic profile, configure the underlying interface on which the router creates the dynamic PPPoE logical interface.

Options ***interface-name***—Variable used to specify the name of the underlying interface on which the PPPoE logical interface is dynamically created. In the **underlying-interface *interface-name*** statement for dynamic PPPoE logical interfaces, you must use the predefined variable **$junos-underlying-interface** in place of ***interface-name***. When the router creates the dynamic PPPoE interface, the **$junos-underlying-interface** predefined variable is dynamically replaced with the name of the underlying interface supplied by the network when the subscriber logs in.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Basic PPPoE Dynamic Profile on page 900
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

unit (Dynamic Demux Interface)

Syntax

```
unit logical-unit-number {
  demux-options {
    underlying-interface interface-name
  }
  family family {
    access-concentrator name;
    address address;
    demux-source {
      source-address;
    }
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
      input filter-name;
      output filter-name;
    }
    mac-validate (loose | strict):
    max-sessions number;
    max-sessions-vsa-ignore;
    rpf-check {
      fail-filter filter-name;
      mode loose;
    }
    service-name-table table-name
    short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
      maximum-seconds>;
    unnumbered-address interface-name <preferred-source-address address>;
  }
  filter {
    input filter-name;
    output filter-name;
  }
}
vlan-id number;
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces demux0]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure a dynamic logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Either the specific unit number of the interface or the unit number variable (**$junos-interface-unit**). The variable is used to specify the unit of the interface when a new demux interface is dynamically created. The static unit number variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

The remaining statements are explained separately.

Copyright © 2013, 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

- Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles on page 771
- For information about static IP demux interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

unit (Dynamic Interface Sets)

Syntax

```
unit logical-unit-number {
  advisory-options {
    downstream-rate rate;
    upstream-rate rate;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces interface-set *interface-set-name* interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Apply the logical interface unit to the interface set.

Options

logical-unit-number—One of the following options:

- **$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **$junos-interface-unit**—For dynamic demux and dynamic PPPoE interfaces, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP or PPP when it accesses the subscriber network.
- *value*—Specific unit number of the interface you want to assign to the dynamic-profile

Range: 0 through 1,073,741,823.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic VLAN Subscriber Interfaces Based on Agent Circuit Identifier Information on page 740
- Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile on page 979
- Configuring an Interface Set of Subscribers in a Dynamic Profile on page 1028
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950

Copyright © 2013, Juniper Networks, Inc.

unit (Dynamic PPPoE)

Syntax

```
unit logical-unit-number {
  keepalives interval seconds;
  no-keepalives;
  pppoe-options {
    underlying-interface interface-name;
    server;
  }
  ppp-options {
    authentication [ authentication-protocols ];
    chap {
      challenge-length minimum minimum-length maximum maximum-length;
    }
    pap;
  }
  family inet {
    unnumbered-address interface-name;
    address address;
    service {
      input {
        service-set service-set-name {
          service-filter filter-name;
        }
        post-service-filter filter-name;
      }
      output {
        service-set service-set-name {
          service-filter filter-name;
        }
      }
    }
    filter {
      input filter-name {
        precedence precedence;
      }
      output filter-name {
        precedence precedence;
      }
    }
  }
  filter {
    input filter-name;
    output filter-name;
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0]

Release Information Statement introduced in Junos OS Release 10.1.

Description In a dynamic profile, configure a logical unit number for the dynamic PPPoE logical interface. You must configure a logical interface to be able to use the router.

Options

logical-unit-number—Variable used to specify the unit number when the PPPoE logical interface is dynamically created. In the **unit *logical-unit-number*** statement for dynamic PPPoE logical interfaces, you must use the predefined variable **$junos-interface-unit** in place of ***logical-unit-number***. The **$junos-interface-unit** predefined variable is dynamically replaced with the unit number supplied by the router when the subscriber logs in.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Basic PPPoE Dynamic Profile on page 900
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

unit (Dynamic Traffic Shaping)

Syntax

```
unit logical-unit-number {
  classifiers {
    type (classifier-name | default);
  }
  output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
  rewrite-rules {
    dscp (rewrite-name | default);
    dscp-ipv6 (rewrite-name | default);
    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
    inet-precedence (rewrite-name | default);
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name*],
[edit dynamic-profiles *profile-name* interfaces interface-set *interface-set-name* interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 9.2.
Support at the **[edit dynamic-profiles *profile-name* class-of-service interfaces interface-set *interface-set-name*]** hierarchy level introduced in Junos OS Release 10.4.

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—One of the following options:

- **$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **$junos-interface-unit**—For dynamic demux and dynamic PPPoE interfaces, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP or PPP when it accesses the subscriber network.
- *value*—Specific unit number of the interface you want to assign to the dynamic-profile

Range: 0 through 16385. For demux and PPPoE interfaces, the unit numbers can range from 0 through 1,073,741,823.

The remaining statements are explained separately. The **classifiers**, **output-traffic-control-profile**, and **rewrite-rules** statements are not supported for interface sets.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile on page 979
- Configuring an Interface Set of Subscribers in a Dynamic Profile on page 1028

Copyright © 2013, Juniper Networks, Inc.

unit (Dynamic Profiles Standard Interface)

Syntax

```
unit logical-unit-number {
  auto-configure {
    agent-circuit-identifier {
      dynamic-profile profile-name;
    }
  }
  dial-options {
    ipsec-interface-id name;
    l2tp-interface-id name;
    (shared | dedicated);
  }
  encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux
    | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap
    | atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr |
    ether-vpls-over-ppp | ethernet | frame-relay-ccc | frame-relay-ppp | frame-relay-tcc |
    frame-relay-ether-type | frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end
    | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc |
    vlan-vci-ccc | vlan-tcc | vlan-vpls);
  family family {
    access-concentrator name;
    address address;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
      adf {
        counter;
        input-precedence precedence;
        not-mandatory;
        output-precedence precedence;
        rule rule-value;
      }
      input filter-name (
        precedence precedence;
      }
      output filter-name {
        precedence precedence;
      }
    }
    max-sessions number;
    max-sessions-vsa-ignore;
    rpf-check {
      fail-filter filter-name;
      mode loose;
    }
    service {
      input {
        service-set service-set-name {
          service-filter filter-name;
        }
        post-service-filter filter-name;
      }
      input-vlan-map {
```

```
            inner-tag-protocol-id tpid;
            inner-vlan-id number;
            (push | swap);
            tag-protocol-id tpid;
            vlan-id number;
          }
          output {
            service-set service-set-name {
              service-filter filter-name;
            }
          }
          output-vlan-map {
            inner-tag-protocol-id tpid;
            inner-vlan-id number;
            (pop | swap);
            tag-protocol-id tpid;
            vlan-id number;
          }
        }
        service-name-table table-name
        short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
          maximum-seconds>;
        unnumbered-address interface-name <preferred-source-address address>;
        filter {
          input filter-name;
          output filter-name;
        }
        keepalives {
          interval seconds;
          }
        ppp-options {
          chap;
          pap;
        }
        vlan-id number;
        vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
      }
    }
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

 Copyright © 2013, Juniper Networks, Inc.

Options *logical-unit-number*—The specific unit number of the interface you want to assign to the dynamic profile, or one of the following Junos OS predefined variables:

- **$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **$junos-interface-unit**—The unit number variable on a dynamic underlying VLAN interface for which you want to enable the creation of dynamic VLAN subscriber interfaces based on agent circuit identifier information.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information on page 737
- Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information on page 739
- Agent Circuit Identifier-Based Dynamic VLANs Components Overview on page 704

unnumbered-address (Dynamic Profiles)

Syntax unnumbered-address *interface-name* <preferred-source-address *address*>;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*]

Release Information Statement introduced in Junos OS Release 9.2.
$junos-preferred-source-address variable support added in Junos OS Release 9.6.
Support for the **$junos-loopback-interface** predefined variable introduced in Junos OS Release 9.6.

Description For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface. To configure unnumbered address dynamically, include the **$junos-loopback-interface-address** predefined variable.

You can configure unnumbered address support on Ethernet interfaces for IPv4 and IPv6 address families.

Options *interface-name*—Name of the interface from which the local address is derived. Use the **$junos-loopback-interface** dynamic variable to dynamically apply a loopback interface. The loopback interface used is based on the routing instance of the subscriber. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.

preferred-source-address *address*—(Optional) Secondary IP address of the donor loopback interface. Use the **$junos-preferred-source-address** dynamic variable to dynamically apply a preferred source address to the unnumbered Ethernet interface. When you use the dynamic variable, the address that is selected resides in the same network as the IP address of the subscriber. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring an Unnumbered Interface* in *Junos OS Network Interfaces Library for Routing Devices*.
- *Junos OS Network Interfaces Library for Routing Devices*

 Copyright © 2013, Juniper Networks, Inc.

unnumbered-address (Dynamic PPPoE)

Syntax unnumbered-address *interface-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" family inet]

Release Information Statement introduced in Junos OS Release 10.1.

Description For dynamic PPPoE interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface.

Options ***interface-name***—Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address.

The **destination** statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Basic PPPoE Dynamic Profile on page 900
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

untagged

Syntax untagged;

Hierarchy Level [edit interfaces ps0]

Release Information Statement introduced in Junos OS Release 13.1.

Description Specify that the router supports untagged traffic on pseudowire subscriber interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Pseudowire Subscriber Logical Interface Device on page 935

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.

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 Authentication and Accounting Parameters for Subscriber Access on page 24

 Copyright © 2013, 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 1339
- Configuring ANCP on page 1332
- Configuring the Method to Set the LAC Connection Speeds to the LNS on page 408

Copyright © 2013, Juniper Networks, Inc.

use-interface-description

Syntax use-interface-description (logical | device);

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* forwarding-options dhcp-relay dhcpv6 group *group-name* relay-agent-interface-id],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* relay-option-82 circuit-id],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-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 group *group-name* relay-agent-interface-id],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-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 group *group-name* relay-option-82 circuit-id],
[edit 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 group *group-name* relay-agent-interface-id],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay relay-option-82 circuit-id],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* relay-option-82 circuit-id]

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.

Description Use the textual interface description instead of the interface identifier when creating the agent-circuit-id suboption of the DHCP relay agent option 82. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

If you specify that the textual description is used and no description is configured for the interface, DHCP relay defaults to using the interface identifier. The textual description is configured using the **description** statement at the **[edit interfaces *interface-name*]** hierarchy level.

NOTE: By default, DHCP relay accepts a maximum of 253 ASCII characters. If the textual interface description is longer than 253 characters, DHCP relay drops the packet, which results in the DHCP client failing to bind.

 Copyright © 2013, Juniper Networks, Inc.

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
- Enabling and Disabling Insertion of Option 82 Information on page 325
- Using a Textual Description in Option 82 on page 327
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 329

use-option-82

Syntax use-option-82;

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 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 13.2.

Description 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.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Enabling and Disabling Insertion of Option 82 Information on page 325
- Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 329

use-primary (DHCP Local Server)

Syntax

use-primary *primary-profile-name*;

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dynamic-profile *profile-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* dynamic-profile *profile-name*],
[edit logical-systems *logical-system-name* system services dhcp-local-server dynamic-profile *profile-name*],
[edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* dynamic-profile *profile-name*],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dynamic-profile *profile-name*],
[edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* dynamic-profile *profile-name*],
[edit system services dhcp-local-server dynamic-profile *profile-name*],
[edit system services dhcp-local-server group *group-name* dynamic-profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the dynamic profile to configure as the primary dynamic profile. The primary dynamic profile is instantiated when the first subscriber or DHCP client logs in. Subsequent subscribers (or clients) are not assigned the primary dynamic profile; instead, they are assigned the dynamic profile specified for the interface. When the first subscriber (or client) logs out, the next subscriber (or client) that logs in is assigned the primary dynamic profile.

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 on page 236

Copyright © 2013, 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 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.

Description

Specify the dynamic profile to configure as the primary dynamic profile. The primary dynamic profile is instantiated when the first subscriber logs in. Subsequent subscribers are not assigned the primary dynamic profile; instead, they are assigned the dynamic profile specified for the interface. When the first subscriber logs out, the next subscriber that logs in is assigned the primary dynamic profile.

Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

EX Series switches do not support DHCPv6.

Options

primary-profile-name—Name of the dynamic profile to configure as the primary dynamic profile

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 236

user-group-profile

Syntax

user-group-profile *profile-name*;

Hierarchy Level

[edit access profile *profile-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Apply a configured PPP group profile to PPP users.

Options

profile-name—Name of a PPP group profile configured at the **[edit access group-profile *profile-name*]** hierarchy level.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Applying a Configured PPP Group Profile to a Tunnel*
- Configuring an L2TP Access Profile on the LNS on page 422

Copyright © 2013, 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.

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 214

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 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.

Description

Specify the user prefix that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

Options

user-prefix-string—User prefix string.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Using External AAA Authentication Services with DHCP on page 214

Copyright © 2013, Juniper Networks, Inc.

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 500
- Configuring the Static Subscriber Global Username on page 503
- Configuring the Static Subscriber Group Username on page 507

Copyright © 2013, Juniper Networks, Inc.

username-include

Syntax

```
username-include {
   circuit-type;
   delimiter delimiter-character;
   domain-name domain-name-string;
   interface-name;
   mac-address;
   option-18;
   option-37;
   option-82 <circuit-id> <remote-id>;
   radius-realm radius-realm-string;
   user-prefix user-prefix-string;
}
```

Hierarchy Level

[edit interfaces *interface-name* auto-configure vlan-ranges authentication],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Configure the username that the router 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 accesses the local authentication service only and does not use external authentication services, such as RADIUS.

The username takes the format ***user-prefix mac-address circuit-type option–82 interface-name domain-name radius-realm*** with each component separated by whatever delimiter you choose.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VLAN Interface Username Information for AAA Authentication on page 729
- Option 82 Suboptions in Authentication Usernames for Autosense VLANs on page 730
- Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs on page 731

Copyright © 2013, Juniper Networks, Inc.

username-include (DHCP Local Server)

Syntax

```
username-include {
  circuit-type;
  client-id;
  delimiter delimiter-character;
  domain-name domain-name-string;
  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 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],
[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 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 statements are explained separately. The option-60 and option-82 statements are not supported in the DHCPv6 hierarchy levels. The client-id, relay-agent-interface-id, relay-agent-remote-id and relay-agent-subscriber-id statements are supported in the DHCPv6 hierarchy levels only.

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 214
- Creating Unique Usernames for DHCP Clients on page 239

username-include (DHCP Relay Agent)

Syntax

```
username-include {
  circuit-type;
  client-id;
  delimiter delimiter-character;
  domain-name domain-name-string;
  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 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.

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:

- **mac-address**
- **option-60**
- **option-82**

The following statements are supported in the DHCPv6 hierarchy levels only:

- **relay-agent-interface-id**
- **relay-agent-remote-id**
- **relay-agent-subscriber-id**

Copyright © 2013, Juniper Networks, Inc.

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

- Creating Unique Usernames for DHCP Clients on page 239
- Using External AAA Authentication Services with DHCP on page 214

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.

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 500
- Configuring the Static Subscriber Global Username on page 503
- Configuring the Static Subscriber Group Username on page 507

valid-lifetime

Syntax valid-lifetime *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix valid. The DHCPv6 server sends this value to the client (router).

If you configure a value for the **valid-lifetime** statement along with a value for the **maximum-lease-time** statement, or if the configured value of the **valid-lifetime** statement is less than the value of the **preferred-lifetime** statement, a commit error occurs. If a value for the **valid-lifetime** statement is not configured, Junos OS uses the same value as the **maximum-lease-time** or **preferred-lifetime** statement. If neither the *maximum-lease-time* nor **preferred-lifetime** statement is configured, Junos OS uses the default value of 86,400 seconds.

Options ***seconds***—Length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix 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 168
- preferred-lifetime on page 1889
- t1-percentage on page 2071
- t2-percentage on page 2072

 Copyright © 2013, 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*
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

variables

Syntax

```
variables {
  variable-name {
    default-value default-value;
    equals expression;
    mandatory;
    radius {
      vendor-id id {
        attribute attribute-number;
        tag tag-number;
      }
    }
    uid;
    uid-reference;
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure user-defined variables in a dynamic profile. The values that the system uses for these variables are applied when the subscriber authenticates.

Options ***variable-name***—Name of the variable.

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

- Dynamic Variables Overview on page 640

 Copyright © 2013, Juniper Networks, Inc.

vendor-id

Syntax vendor-id *id*;

Hierarchy Level [edit dynamic-profiles *profile-name* variables radius]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the vendor ID as a variable in a dynamic profile.

Options *id*—Vendor ID for the RADIUS attribute.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring User-Defined CoS Variables in a Dynamic Service Profile on page 1000

vendor-specific-tags (Dynamic Traffic Shaping)

Syntax
```
vendor-specific-tags actual-data-rate-downstream;
vendor-specific-tags access-loop-encapsulation;
```

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]

Release Information Statement introduced in Junos OS Release 12.1.

Description Set the shaping-rate and overhead-accounting class-of-service attributes based on Vendor-Specific Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101].

Options **vendor-specific-tags** can be set to one or both of the following:

- **access-loop-encapsulation**—Set the overhead-accounting class-of-service attribute based on access line parameters in PPPoE discovery packets on dynamic subscriber interfaces.
- **actual-data-rate-downstream**—Set the shaping-rate class-of-service attribute based on the actual-data-rate-downstream attribute.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 1074
- Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces on page 1093
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 1072

Copyright © 2013, Juniper Networks, Inc.

version (BFD)

Syntax version (0 | 1 | automatic);

Hierarchy Level [edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the BFD protocol version to detect.

Options **0**—Use BFD protocol version 0.

1—Use BFD protocol version 1.

automatic—Autodetect the BFD protocol version.

Default: automatic

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 269
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 357

 Copyright © 2013, Juniper Networks, Inc.

version (Dynamic IGMP Interface)

Syntax version *version*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols igmpinterface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the version of IGMP.

Options *version*—IGMP version number.

Range: 1, 2, or 3

Default: IGMP version 2

NOTE: Routers running different versions of IGMP negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.

If you have already configured the router to use IGMP version 1 and then configure it to use IGMP version 2, the router continues to use IGMP version 1 for up to 6 minutes and then uses IGMP version 2.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Client Access on page 679
- For information about specifying a different IGMP version, see “Changing the IGMP Version” in the *Multicast Protocols Feature Guide for Routing Devices*

version (Dynamic MLD Interface)

Syntax version *version*;

Hierarchy Level [edit dynamic-profiles *profile-name* protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the MLD version explicitly on the dynamic interface. MLD version 2 (MLDv2) is used only to support source-specific multicast (SSM).

Options ***version***—MLD version to run on the interface.

Range: 1 or 2

Default: 1 (MLDv1)

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- *Modifying the MLD Version*

 Copyright © 2013, Juniper Networks, Inc.

virtual-network

Syntax

```
virtual-network {
  home-agent-address ip-address {
    registration-lifetime seconds;
    revocation-required;
    timestamp-tolerance seconds;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip home-agent],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent],
[edit services mobile-ip home-agent]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the virtual network for the Mobile IP home agent. Only one virtual network is supported.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Mobile IP Home Agent on page 570

vlan-id (Dynamic Profiles)

Syntax vlan-id (*number* | none);

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.
VLAN demux interface support introduced in Junos OS Release 10.2.

Description For VLAN demux, Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.

Options *number*—A valid VLAN identifier. When used in the **dynamic-profiles** hierarchy, specify the **$junos-vlan-id** predefined variable to dynamically obtain the VLAN identifier.

none—Enable the use of untagged pseudo-wire frames on dynamic interfaces.

Range:

- For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.
- For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.
- VLAN ID 0 is reserved for tagging the priority of frames.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Static Subscriber Interfaces Using VLAN Demux Interfaces on page 767
- Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles on page 772

 Copyright © 2013, Juniper Networks, Inc.

vlan-id (Dynamic VLANs)

Syntax vlan-id *number*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 10.4.

Description For dynamic VLAN interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.

You cannot include the **vlan-id** statement with the **swap** statement, **swap-push** statement, **push-push** statement, or **push-swap** statement at the **[edit** dynamic-profiles ***profile-name*** interfaces ***interface-name*** unit ***logical-unit-number*** output-vlan-map**]** hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the **vlan-id** statement that you include at the **[edit** dynamic-profiles ***profile-name*** interfaces ***interface-name*** unit ***logical-unit-number***] hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Rewriting the VLAN Tag on Tagged Frames*
- *Binding VLAN IDs to Logical Interfaces*

vlan-model

Syntax vlan-model one-to-one;

Hierarchy Level [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 11.2.

Description Define the network VLAN model.

Options **one-to-one**—Specify that any received, dual-tagged VLAN packet triggers the provisioning process in a Layer 2 Wholesale network. Using this option, the router learns VLAN tags for each individual client. The router learns both the outer tag and inner tag of the incoming packets, when the **instance-role** statement is defined as **access**, or the outer VLAN tag only, when the **instance-role** statement is defined as **nni**.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers*
- *Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers*
- *Junos OS Subscriber Management and Services Library*

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 47
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24

 Copyright © 2013, 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 66
- Guidelines for Configuring RADIUS NAS-Port Options for Subscriber Access per Physical Interface, VLAN, or Stacked VLAN on page 65

Copyright © 2013, Juniper Networks, Inc.

vlan-tag (Dynamic Classifiers)

Syntax vlan-tag (inner | outer);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers ieee-802.1]

Release Information Statement introduced in Junos OS Release 10.1.

Description Apply this IEEE-802.1 classifier to the inner or outer VLAN tags in a dynamic profile.

Default If you do not include this statement, the classifier applies to the outer VLAN tag only.

Options **inner**—Apply the classifier to the inner VLAN tag only.

outer—Apply the classifier to the outer VLAN tag only.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying a Classifier to a Subscriber Interface in a Dynamic Profile on page 981
- *classifiers (Definition)*

 Copyright © 2013, Juniper Networks, Inc.

vlan-tag (Dynamic Rewrite Rules)

Syntax vlan-tag (outer | outer-and-inner);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules ieee-802.1]

Release Information Statement introduced in Junos OS Release 10.1.

Description Apply this IEEE-802.1 rewrite rule to the outer or outer and inner VLAN tags in a dynamic profile.

Default If you do not include this statement, the rewrite rule applies to the outer VLAN tag only.

Options **outer**—Apply the rewrite rule to the outer VLAN tag only.

outer-and-inner—Apply the rewrite rule to both the outer and inner VLAN tags.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 950
- Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile on page 980
- *rewrite-rules*

vlan-tagging

Syntax vlan-tagging;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name*],
[edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description For Fast Ethernet and Gigabit Ethernet interfaces and aggregated Ethernet interfaces configured for VPLS, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

NOTE: For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or bridging domain.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VLAN Ranges for Use with Dynamic Profiles on page 724
- Configuring Static Subscriber Interfaces in Dynamic Profiles on page 765
- Configuring the L2TP LNS Peer Interface on page 426

 Copyright © 2013, Juniper Networks, Inc.

vlan-tags

Syntax vlan-tags outer [*tpid*].*vlan-id* [inner [*tpid*].*vlan-id*];

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.
VLAN demux interface support introduced in Junos OS Release 10.2.

Description For Gigabit Ethernet IQ and IQE interfaces only, binds TPIDs and 802.1Q VLAN tag IDs to a logical interface. You must include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level.

NOTE: The **inner-range *vid1–vid2*** option is supported on MX Series routers with IQE PICs only.

Options

inner [*tpid*].*vlan-id*—A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the **$junos-vlan-id** predefined variable to dynamically obtain the VLAN ID.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer [*tpid*].*vlan-id*—A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the **$junos-stacked-vlan-id** predefined variable.

Range: For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Dual VLAN Tags*
- *stacked-vlan-tagging*

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 28
- Configuring Authentication and Accounting Parameters for Subscriber Access on page 24

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.

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 36
- Configuring Per-Subscriber Session Accounting on page 29

 Copyright © 2013, Juniper Networks, Inc.

weighted-load-balancing (L2TP LAC)

Syntax weighted-load-balancing;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify that the router chooses among multiple tunnels that share the same preference level by considering the maximum sessions configured per tunnel. The tunnel configured with the highest maximum number of sessions in the preference level has the highest weight. This tunnel is selected until the maximum number of sessions for the tunnel is reached. Then the router selects the tunnel with the next higher weight to establish connections until that tunnel's maximum session limit is reached, and so on.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Weighted Load Balancing for LAC Tunnel Sessions on page 407
- Configuring the L2TP LAC Tunnel Selection Parameters on page 406

wimax

Syntax wimax;

Hierarchy Level [edit services mobile-ip access-type],
[edit logical-systems *logical-system-name* services mobile-ip access-type],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* services mobile-ip access-type],
[edit routing-instances *routing-instance-name* services mobile-ip access-type]

Release Information Statement introduced in Junos OS Release 9.5.

Description Enable WiMAX features for Mobile IP home agent, including the ability to process, send, and receive WiMAX Vendor Specific Attributes (VSAs).

NOTE: Although this statement is available in the CLI for both default and nondefault router contexts, the commit operation is disallowed when you configure the statement in a nondefault router context.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 569
- Configuring the Access Type for Mobile IP on page 573

wins-server (Access)

Syntax

```
wins-server {
   ipv4-address;
}
```

Hierarchy Level

[edit access address-assignment pool *pool-name* family inet dhcp-attributes]

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 168

Copyright © 2013, Juniper Networks, Inc.

PART 20

Index

- Index on page 2123
- Index of Statements and Commands on page 2167

Copyright © 2013, Juniper Networks, Inc.

Index

Symbols

#, comments in configuration statements.....lxxii
$junos-cos-byte-adjust predefined
variable.....1488, 1494, 1651, 1841
$junos-cos-byte-adjust-cell predefined
variable.....1494, 1841
$junos-cos-byte-adjust-frame predefined
variable.....1651, 1841
$junos-cos-delay-buffer-rate predefined
variable.....1524
$junos-cos-excess-priority predefined
variable.....1620
$junos-cos-excess-rate predefined variable.....1622
$junos-cos-excess-rate-high predefined
variable.....1623
$junos-cos-excess-rate-low predefined
variable.....1624
$junos-cos-guaranteed-rate predefined
variable.....1672
$junos-cos-overhead-accounting predefined
variable.....1841
$junos-cos-scheduler predefined
variable.....989, 1970
$junos-cos-scheduler-bs predefined
variable.....989, 1483
$junos-cos-scheduler-dropfile-any predefined
variable.....989, 1580
$junos-cos-scheduler-dropfile-high predefined
variable.....989, 1580
$junos-cos-scheduler-dropfile-low predefined
variable.....989, 1580
$junos-cos-scheduler-dropfile-medium-high
predefined variable.....989, 1580
$junos-cos-scheduler-dropfile-medium-low
predefined variable.....989, 1580
$junos-cos-scheduler-excess-rate predefined
variable.....1621
$junos-cos-scheduler-map predefined
variable.....1968
$junos-cos-scheduler-pri predefined
variable.....989, 1896
$junos-cos-scheduler-shaping-rate predefined
variable.....1988
$junos-cos-scheduler-tx predefined
variable.....989, 2057
$junos-cos-shaping-mode predefined
variable.....1651, 1841
$junos-cos-shaping-rate predefined
variable.....1988
$junos-subscriber-ipv6-multi-address predefined
variable.....231
( ), in syntax descriptions.....lxxii
3GPP Gx
Gx-Plus and.....541
802.1Q VLANs
dynamic.....1471
VLAN tagging.....2116
< >, in syntax descriptions.....lxxi
[ ], in configuration statements.....lxxii
{ }, in configuration statements.....lxxii
| (pipe), in syntax descriptions.....lxxii

A

AAA
ANCP interactions.....1317
configuration testing.....156
Mobile IP home agent and 560
with Diameter base protocol.....455
AAA access messages
supported attributes.....110
AAA accounting messages
supported attributes.....115
AAA directed logout
DHCP authentication services.....214, 288
AAA logical system/routing instance
domain map.....184
AAA Service Framework.....21
dynamic service activation
during login.....86
aaa-access-profile statement
L2TP LNS.....1413
aaa-logical-system statement
domain map.....1414
aaa-routing-instance statement
domain map.....1415
abated-utilization statement
address-assignment pools.....1415
abated-utilization-v6 statement
address-assignment pools.....1416

access loop circuit identifier
ANCP....1313
configuring ANCP....1334
access profile
configuring global static subscriber....501
configuring static subscriber group....505
domain map....182
session options....153
access profile statements
access-profile-name....1424
accounting-backup-options....1428
duplication-vrf....1587
max-pending-accounting-stops....1780
max-withhold-time....1785
session-options....1987
update-interval....2088
vrf-name....2118
access profiles
attaching....150
configuring....149
access routes
configuring....690
delaying removal for DHCP and PPP subscribers....693, 1657
removal for DHCP and PPP subscribers....688
subscriber management....687
verifying....694
access statement
dynamic profiles....1417
access type
configuring Mobile IP....573
access-concentrator statement....1418
access-identifier statement
ANCP....1419
access-internal routes
configuring for DHCP....692
configuring for PPP....692
delaying removal for DHCP and PPP subscribers....693, 1657
removal for DHCP and PPP subscribers....688
subscriber management....687
verifying....694
access-internal statement
dynamic profiles....1420
access-loop-id-local statement....1421
access-profile statement
domain map....1421
routing instances....1422
static subscribers....1423
access-profile statements
accounting....1426
access-profile-name statement
duplicate accounting....1424
access-type statement
Mobile IP....1425
accounting
configuring RADIUS....24
duplicate reports....28
Mobile IP time-based....563
opaque DHCP options....83
accounting methods....24
configuring Mobile IP....572
accounting statement
access profile....1426
dynamic IGMP interface....1427
dynamic MLD interface....1427
accounting statistics....26
per-service accounting....26
subscriber service session....35, 1175
subscriber session....29
accounting-backup-options statement
RADIUS accounting....1428
accounting-order statement
service accounting....1428
accounting-port statement....1429
accounting-server statement....1429
accounting-session-id-format statement....1430, 1739
accounting-stop-on-access-deny statement....1430
accounting-stop-on-failure statement....1431
Acct-Off messages....25
Acct-On messages....25
ACI
ANCP....1313
ACI (agent circuit identifier) interface sets
adjusting CoS shaping rate and overhead accounting....742
applying CoS traffic shaping attributes....996, 998, 1003
associating with subscriber interfaces....1727
bandwidth management overview....706
bandwidth management restrictions....709
clearing....748
configuration overview....733
configuring....1446, 1472, 1594

Copyright © 2013, Juniper Networks, Inc.

configuring dynamic subscriber interfaces based on ACI information....740
configuring dynamic underlying interfaces....737
configuring PPPoE options....1879
configuring static underlying interfaces....739
CoS traffic shaping predefined variables....1005
defining....735, 1710, 1728
overview....702
overview of configuration components....704
verifying and managing ACI-based dynamic VLANs....747
action statement....1432
activating subscriber services
CLI-based....1122
active server groups
DHCP relay....328
active-server-group statement
DHCP relay agent....1433
address assignment precedence....44
address pool
domain map....184
address statement
Diameter base protocol
peer....1434
transport....1435
interface....1434
tunnels
LAC....1436
LNS....1435
address-assignment pool
configuring for L2TP LNS....424
address-assignment pools
abated-utilization....1415
abated-utilization-v6....1416
client attributes....172
configuring....168
DHCP attributes....173
mapping option 82....173
setting the grace period....173
setting the maximum lease time....173
setting the name server address....173
specifying NetBIOS node type....173
specifying router addresses....173
specifying the boot file 173
specifying the boot server 173
specifying the DNS server IPv6 address....173
specifying the domain name to search....173
specifying the SIP server domain name....173
specifying the SIP server IPv6 address....173
specifying the source address....173
specifying the TFTP server....173
specifying the WINS server....173
specifying user-defined options....173
DHCP local server....202
DHCPv6 attributes....173
specifying the preferred-lifetime....173, 1542
specifying the t1-percentage....173, 1542
specifying the t2-percentage....173, 1542
specifying the valid-lifetime....173, 1542
high-utilization....1678
high-utilization-v6....1679
license requirements....175
linking....170
name....169
named range....169
network address....169
router advertisement....1811
static address....171
threshold traps....172
tracing operations....175
address-assignment statement
address-assignment pools....1437
address-change-immediate-update statement
accounting....1438
address-pool statement
domain map....1438
address-protection statement
access....1439
adf statement
dynamic firewalls....1440
adjacency timer
ANCP global configuration....1335
ANCP neighbor configuration....1333, 1335
adjacency-timer statement
ANCP....1441
adjust-minimum statement....1443
adjust-percent statement....1444
adjustment-control-profiles
configuring....1106
effective shaping rate....1084

overview....1084
viewing....1107
adjustment-control-profiles statement
CoS....1442
advisory-options statement
ANCP....1445
L2TP....1445
Agent Circuit ID suboption
DHCP relay....325
agent circuit identifier *See* ACI
agent-circuit-identifier statement
dynamic VLAN interface sets....1446
aggregate (logical interface) policer
configuration statement for....1759
aggregate statement
hierarchical policer....1447
aggregate-clients statement
DHCP local server....1448
DHCP relay agent....1449
static subscribers....1451
aggregated Ethernet
targeted distribution of
subscribers....1758, 1926, 2015
aggregated Ethernet interfaces without link
protection
manually rebalancing subscribers 830
periodically rebalancing subscribers 830
rebalancing subscribers....830
targeting subscribers....830
aggregated Ethernet logical interfaces *See*
subscriber interfaces, IP demux over aggregated
Ethernet *See* subscriber interfaces, VLAN demux
over aggregated Ethernet *See* subscriber
interfaces, VLAN over aggregated Ethernet
algorithm statement
Mobile IP....1452
allow-snooped-clients statement
DHCP relay agent....1453
always-write-giaddr statement
DHCP relay agent....1454
always-write-option-82 statement
DHCP relay agent....1455
anchor logical tunnel....930
anchor-point statement
pseudowires....1456
ANCP
AAA interactions....1317
access loop circuit identifier....1313
access loop circuit identifier
configuration....1334
ACI....1313
adjacency timer global configuration....1335
adjacency timer neighbor configuration....1335
adjusting subscriber line traffic
rates....1337, 1339, 1341
adjusting subscriber traffic with....1314
associating subscriber VLANs with access
nodes....1334
backwards compatibility global
configuration....1336
backwards compatibility neighbor
configuration....1333, 1336
configuration overview....1332
CoS information
verifying....1345
CoS shaping across restarts....1329
CoS shaping rate adjustment for subscriber
local loops....1314
discovery table global configuration....1335
discovery table neighbor
configuration....1333, 1335
event logging....1345
flags for tracing operations....1348
log file access for tracing operations....1347
log file size and number....1346
log filenames for tracing operations....1346
message severity levels for tracing
operations....1348
monitoring subscriber traffic with....1311
neighbor configuration
adjacency timer....1333
discovery table entries....1333
ietf mode....1333
IP address....1333
pre-ietf mode....1333
neighbor information
clearing....1343
verifying....1343
operations in different access models....1319
overview....1311
pre-ietf mode global configuration....1336
pre-ietf mode neighbor
configuration....1333, 1336
regular expressions for tracing
operations....1347
reporting traffic to CoS....1338
restart time global configuration....1337

Copyright © 2013, Juniper Networks, Inc.

shaping-rate adjustments for subscriber local loops....1082
subscriber information
clearing....1344
verifying....1344
tracing operations....1345
triggering OAM....1342
ANCP agent *See* ANCP
ancp statement
ANCP....1457
CoS....1442, 1458, 1459
ANCP statements
access-identifier....1419
adjacency-timer....1441
advisory-options....1445
ancp....1457
downstream-rate....1577
ietf-mode....1690
interface-set....1725
interfaces....1732
juniper-dsl-attributes....1744
maximum-discovery-table-entries....1777
maximum-helper-restart-time....1777
neighbor
for all neighbors....1811
for unique access loop circuit identifier....1810
pre-ietf-mode....1893
qos-adjust....1910
qos-adjust-adsl....1910
qos-adjust-adsl2....1911
qos-adjust-adsl2-plus....1911
qos-adjust-sds1....1912
qos-adjust-vds1....1912
qos-adjust-vds2....1913
traceoptions....2033
underlying-interface....2073
upstream-rate....2089
application statement....1459
CoS....1442, 1459
apply-groups statement
subscriber secure policy....1460
apply-groups-except statement
subscriber secure policy....1460
Ascend-Data-Filter
example of dynamic configuration....1202
example of static configuration....1205
field descriptions....1138
multiple filters....1137
naming convention....1137
verifying configuration....1179
Ascend-Data-Filters....1136
Ascent-Data-Filter....1177
assignment-id-format statement....1461
Asynchronous Transfer Mode. *See* ATM
ATM (Asynchronous Transfer Mode)
configuration examples
dynamic PPPoE subscriber interface over ATM....593
static IP subscriber interface over ATM....610
static IP subscriber interface over Ethernet over ATM....616
static PPP subscriber interface over ATM....623
static PPPoE subscriber interface over ATM....602
configuration overview for subscriber access....587
encapsulation types for subscriber access....585
subscriber access configuration guidelines....589
subscriber access overview....581
verifying and managing subscriber access configurations....590
ATM subscriber access
configuration examples
dynamic PPPoE over ATM....593
static IP over ATM....610
static IP over Ethernet over ATM....616
static PPP over ATM....623
static PPPoE over ATM....602
configuration guidelines....589
configuration overview....587
dynamic PPPoE access over ATM network....593
encapsulation types....585
overview....581
static PPPoE access over ATM network....602
verifying and managing....590
attempts statement
DHCP local server....1462
attribute statement
dynamic profile variables....1463
attributes statement....1463
authenticate statement
Mobile IP....1464

Copyright © 2013, Juniper Networks, Inc.

authentication
configuring RADIUS....24
dynamic VLAN....728
Mobile IP home agent....560
authentication and accounting information
retaining....31
authentication attributes
local Mobile IP....571
authentication method
configuring Mobile IP....570
authentication methods....24
authentication password *See* password
configuring global static subscriber....503
configuring group static subscriber....506
authentication protocols
controlling order for PPP....367, 1467
modifying the length of the CHAP
challenge....369, 1495
authentication services
with DHCP....214, 288
authentication statement
DHCP local server....1465
DHCP relay agent....1466
dynamic PPP....1467
static subscribers....1468
authentication-order statement
access....1469
authentication-server statement....1470
authorization, subscriber
JSRC....493
authorization-order statement
JSRC....1470
auto logout
DHCP....225, 308
DHCP relay agent option 82....310
auto-configure
dynamic VLAN interface sets....1472
auto-configure statement....1471
autonomous statement
dynamic router advertisement....1473
avp statement
L2TP tunnel switch profiles....1473
AVPs
Diameter....462
Juniper Networks
Diameter and Diameter applications....462

B

backup-options, RADIUS accounting
configuring....38, 39
overview....32, 39
backwards compatibility
ANCP global configuration....1336
ANCP neighbor configuration....1333, 1336
bandwidth management
agent circuit identifier-based dynamic
VLANs....706, 709
configuring....742
bandwidth statement
inline services....1474
bandwidth-limit statement
hierarchical policer....1475
policer....1476
bandwidth-percent statement
policer....1478
bearer-type statement
L2TP tunnel switch profiles....1480
bfd statement
liveness detection method....1481
bindingst
clearing DHCP....234, 315
boot-file statement....1482
boot-server statement....1482
BOOTREPLY packets
DHCP snooping....299
braces, in configuration statements....lxxii
brackets
angle, in syntax descriptions....lxxi
square, in configuration statements....lxxii
buffer-size statement
dynamic CoS....1483
burst-size
configuring for MIC and MPC interfaces....1070
burst-size-limit statement....1485
hierarchical policer....1484
bytes statement
CoS statements....1488, 1494, 1651
dynamic CoS statements....1488, 1494

C

Calling Number AVP 22
preventing L2TP LAC from sending....407
calling-number statement
L2TP tunnel switch profiles....1489
Calling-Station-ID
configuring....56

Copyright © 2013, Juniper Networks, Inc.

calling-station-id-delimiter statement.....1490
calling-station-id-format statement.....1491
captive portal content delivery
    dynamic subscriber interfaces.....1229
captive-portal-content-delivery statement.....1492
captive-portal-content-delivery statements
    traceoptions.....2035
captive-portal-content-delivery-rule-set
  statement.....1493
cell-mode statement
    CoS statements.....1494
    dynamic CoS statements.....1494
challenge-length statement
    dynamic PPP.....1495
change of authorization *See* CoA
CHAP challenge
    modifying length of.....369
chap statement.....1496, 1497
    dynamic PPP.....1497
chassis statements
    enhanced-policer.....1615
circuit-id statement
    address-assignment pools.....1498
    DHCP relay agent.....1499
circuit-type statement
    DHCP local server.....1500
    DHCP relay agent.....1501
cisco-nas-port-info statement
    L2TP tunnel switch profiles.....1502
class of service.....949, 956 *See* CoS
    See also CoS
class-of-service statement
    subscriber access.....1503
classic filters
    components.....1133
    processing order.....1134
    types.....1133
classic firewall filters
    configuration guidelines.....1135, 1164
classifiers statement
    dynamic CoS.....1503
clear-on-abort statement
    DHCP local server.....1504
client attributes
    address-assignment pools.....172
client configuration information
    DHCP.....201
client sessions.....1164
client statement.....1505
client usernames
    DHCP
        unique.....239
client-accounting-algorithm statement
    RADIUS.....1506
client-authentication-algorithm statement
    RADIUS.....1507
client-discover-match statement
    DHCP local server.....1508
    DHCP relay agent.....1509
client-id statement.....1510
    DHCP relay agent.....1511
client-idle-timeout statement
    access profile session options.....1511
client-session-timeout statement
    access profile session options.....1512
CoA
    messages.....87
    RADIUS.....87
coa-dynamic-variable-validation.....1512
coa-immediate-update statement
    accounting.....1513
color-aware statement.....1514
color-blind statement.....1515
comments, in configuration statements.....lxxii
committed-burst-size statement.....1516
committed-information-rate statement.....1518
concurrent-data-sessions statement.....1519
connect-actively statement
    Diameter base protocol.....1520
conventions
    text and syntax.....lxxi
CoS
    hierarchical scheduling
        MX Series routers.....956
        three-level.....956, 1677
        two-level.....956, 1677
    implicit-hierarchy.....956
    implicit-hierarchy statement.....1677
    maximum–hierarchy–levels.....956
    maximum–hierarchy–levels statement.....1677
    RADIUS-provided parameters
        configuring an access dynamic
          profile.....999
        example.....1040
        overview.....989
    reporting effective shaping rate.....1108
    shaping preserved across ANCP restarts.....1329

Copyright © 2013, Juniper Networks, Inc.

shaping-rate adjustments for agent circuit identifier-based dynamic VLANs.....706, 709
configuring.....742
shaping-rate adjustments for subscriber local loops
configuration guidelines.....1083
disabling.....1104
enabling.....1099
example.....1114
overview.....1082
subscriber access
changing services.....993
classifiers.....981
configuration guidelines.....950
configuration overview.....963
configuring variables.....1000
dynamic configuration overview.....965
interfaces.....979
modifying.....1125
overview.....949
rewrite rules.....980
static scheduling and queuing example.....1031
traffic parameters.....971, 972
traffic-control profiles.....1125
CoS statements
adjustment-control-profiles.....1442
ancp.....1442, 1458, 1459
application.....1442, 1459
bytes statement.....1488, 1494, 1651
cell-mode statement.....1494
effective-shaping-rate.....1605
frame-mode statement.....1651
pppoe-tags.....1442, 1459, 1878
radius-coa.....1442, 1459, 1918
CoS traffic shaping
with ANCP.....1339, 1341
curly braces, in configuration statements.....lxxii
current-hop-limit statement
dynamic router advertisement.....1520
customer support.....lxxii
contacting JTAC.....lxxii

D

database-replication statement
subscriber session database replication.....1521
deactivating subscriber services
CLI-based.....1122
default subscriber service.....1127
configuring.....1127
default-action statement
DHCP relay agent.....1522
default-lifetime statement
dynamic router advertisement.....1523
default-value statement
dynamic profile variables.....1523
delay-buffer-rate statement
dynamic CoS.....1524
delegated-pool statement
DHCP local server.....1525
delimiter statement
DHCP local server.....1526
DHCP relay agent.....1528
domain map.....1529
delimiters
domain names.....187
demux interfaces
unit statement.....2076
demux logical interfaces *See* subscriber interfaces, IP demux over VLAN demux
demux-options statement
dynamic IP demux interface.....1531
demux-source statement
dynamic IP demux interfaces.....1532
dynamic underlying interface.....1533
demux0 statement
dynamic IP demux interface.....1530
demux0 statements
underlying-interface.....2074
destination lockout timeout
L2TP.....414
destination statement
Diameter base protocol.....1533
L2TP LNS.....1534
destination-address statement
cpcd.....1534
subscriber secure policy.....1535
destination-host statement
Gx-Plus.....1536
JSRC.....1535
PTSP.....1536
destination-port statement
subscriber secure policy.....1537, 1583
destination-prefix-list statement
captive-portal-content-delivery.....1537

Copyright © 2013, Juniper Networks, Inc.

destination-realm statement
Gx-Plus....1538
JSRC....1538
PTSP....1539
destruct-timeout
L2TP tunnels....413
destruct-timeout statement....1539
detection-time statement
liveness detection....1540
device-count statement
pseudowires....1541
DHCP
authentication services....214, 288
AAA directed logout....214, 288
auto logout....225, 308
centrally-configured DHCP options....80, 84
clearing bindings....234, 315
client configuration information....201
default route installation
preventing....222, 223, 297, 298
delaying removal of access and access-internal routes....693, 1657
distinguishing duplicate clients....214, 288
duplicate client IDs....212, 286
duplicate clients
configuration guidelines....213, 287
duplicate MAC addresses....212, 286
firewall filters on routers using the jdhcpd process....208
example....265
grouping interfaces....217, 290
configuration guidelines....218, 291
information request processing
overriding....227
maximum clients per interface
overriding....221, 298
message severity levels for tracing operations....259, 337
override settings
deleting....229
rapid commit....228
regular expressions for tracing operations....258, 336
removal of access and access-internal routes....688
unified ISSU....7
unique client usernames....239
user passwords....238
DHCP local server
address-assignment pool selection....215
address-assignment pools....202
client auto logout....226, 311
default route installation
preventing....222, 223, 297, 298
delegated IPv6 pool....231
DHCP snooping....223, 299
DHCPv6....204
DHCPv6 rapid commit....228
dynamic client reconfiguration
authentication token configuration....248
behavior on failure configuration....247
configuration overview....245
number of attempts configuration....246
preventing binding of nonsupporting clients....249
RADIUS-initiated disconnect configuration....247
requesting....248
dynamic profile attachment
multiple subscribers....236, 330
overview....207, 282
use primary profile....236, 330
graceful Routing Engine switchover....254
grouping interfaces
options....219
information request message processing....205
information request processing
overriding....227
interaction
address-assignment pools....200
DHCP clients....200
liveness detection....203, 251
maximum clients per interface
overriding....221, 298
minimal configuration
default settings....202
option 50....230
override settings
deleting....229
overriding default configuration....220
overview....200
per-interface tracing operations....260, 338
processing client information requests....205
specific address....230
tracing operations....255, 333
verifying configuration....253

Copyright © 2013, Juniper Networks, Inc.

DHCP local server statements
attempts....1462
boot-file....1482
boot-server....1482
circuit-type....1500
clear-on-abort....1504
client-discover-match....1508
client-id....1510
delegated-pool 1525
delimiter....1526
dhcp-local-server....1543
dhcpv6....1554
dns-server....1567
domain-name....1570
duplicate-clients-on-interface....1585
dynamic-profile....1590
forward-snooped-clients....1645
group....1658
interface....1705
interface-client-limit....1717
interface-name....1723
interface-traceoptions....1730
ip-address-first....1740
logical-system-name....1762
mac-address....1767
option-60....1826
option-82....1829, 1830
overrides....1842
password....1854
pool....1870
pool-match-order....1872
process-inform....1897
rapid-commit....1925
relay-agent-interface-id....1933
relay-agent-remote-id....1934
relay-agent-subscriber-id....1935
route-suppression....1953
routing-instance-name....1958
service-profile....1982
strict....2006
timeout....2024
trace....2029
traceoptions....2037
trigger....2059
username-include....2099
DHCP option information
DHCP relay 321, 323
DHCP option strings
DHCP relay agent....349, 353
DHCPv6 relay agent....349, 353
DHCP options
accounting....83
configuring on RADIUS....80
opaque....80
RADIUS-sourced....80
renewing....83
verifying....84
DHCP relay
access and access-internal routes....687
active server groups....328
Agent Circuit ID suboption....325
automatic binding of stray requests....320
client auto logout....226, 311
configuration examples
minimum configuration....341
multiple clients and servers
configuration....342
default route installation
preventing....222, 223, 297, 298
delaying removal of access and access-internal
routes....693, 1657
DHCP option information....321, 323
DHCP relay proxy....280, 329
DHCP snooping....299, 301, 345
DHCPv6....279
DHCPv6 Interface-ID option....329
disabling....319, 339
discarded packets
counting....325
dynamic profile attachment
multiple subscribers....236, 330
overview....207, 282
use primary profile....236, 330
graceful Routing Engine switchover....254
grouping interfaces
options....292
how components interact....278
Layer 2 unicast transmission....296
liveness detection....279, 318
maximum clients per interface
overriding....221, 298
option 82
auto logout....310
option 82 information....325
option 82 prefix....326
option 82 textual description....327

Copyright © 2013, Juniper Networks, Inc.

override settings
deleting....229
overriding broadcast bit....296
overriding default configuration....293
overriding option 82....295
overview....276
overwrite giaddr....294
per-interface tracing operations....260, 338
removal of access and access-internal routes....688
replacing IP source address....295
selective traffic processing....321, 323, 349, 353
sending release messages....312
server groups....328
tracing operations....255, 333
trusting option 82....296
verifying configuration....332
DHCP relay agent
DHCP option strings....349, 353
DHCP snooping....306
DHCP relay agent statements
active-server-group....1433
allow-snooped-clients....1453
always-write-giaddr....1454
always-write-option-82....1455
circuit-id....1499
circuit-type....1501
client-discover-match....1509
client-id....1511
default-action....1522
delimiter....1528
dhcp-relay....1548
dhcpv6....1557
disable-relay....1567
domain-name....1572
drop....1578
duplicate-clients-on-interface....1585
dynamic-profile....1591
equals....1616
forward-only....1644
forward-snooped-clients....1646
group....1661
interface....1707
interface-client-limit....1719
interface-name....1724
interface-traceoptions....1730
layer2-unicast-replies....1752
local-server-group....1757
logical-system-name....1763
mac-address....1768
no-allow-snooped-clients....1817
no-bind-on-request....1818
option-60....1827
option-number....1833
overrides....1844
password....1855
prefix....1891
proxy-mode....1907
relay-agent-interface-id....1930
username....1931
relay-agent-remote-id....1932
relay-agent-subscriber-id....1936
relay-option....1937
relay-option-82....1938
relay-server-group....1939
replace-ip-source-with....1941
route-suppression....1953
routing-instance-name....1959
send-release-on-delete....1973
server-group....1975
service-profile....1983
starts-with....2001
trace....2030
traceoptions....2037
trust-option-82....2060
use-interface-description....2090
use-option-82....2091
use-primary....2092, 2093
user-prefix....2095, 2096
username-include....1466, 2100
DHCP relay proxy....1907
enabling....329
how components interact....281
overview....280
DHCP snooping
BOOTREPLY packets....299
DHCP local server....223, 299
DHCP relay agent....306
disabling....301
disabling interfaces....299
enabling....301
enabling for DHCPv6 relay agent....345
enabling interfaces....299
example of DHCP relay agent configuration....269, 343, 357, 1285
example of enabling for DHCPv6 relay agent....345

Copyright © 2013, Juniper Networks, Inc.

DHCP stray requests
disabling automatic binding....320
enabling automatic binding....320
DHCP subscriber
auto logout....226, 311
liveness detection....251, 318
dhcp-attributes statement
address-assignment pools....1542
dhcp-local-server statement....1543
dhcp-relay statement....1548
DHCPv6
multiple address assignment....230
DHCPv6 Interface-ID option
DHCP relay....329
DHCPv6 local server
IA_NA....230
overview....204
specific address....230
verifying configuration....253
DHCPv6 local server statements
preferred-lifetime....1889
t1-percentage....2071
t2-percentage....2072
valid-lifetime....2102
DHCPv6 Option-18
DHCP relay....329
DHCPv6 relay
DHCP snooping....301, 345
overview....279
verifying configuration....333
DHCPv6 relay agent
DHCP option strings....349, 353
dhcpv6 statement....1554, 1557
dial-options statement....1560
dynamic profiles....1561
Diameter
AVPs....462
message sequences for Gx-Plus....543
message sequences for JSRC....488
message sequences for PTSP....517
messages used by Diameter applications....457
Diameter base protocol....455
configuration overview....471
event log access....477
event logging....475
filtering trace operation output....477
flags for tracing operations....478
function information
verifying....479
instance information
verifying....478
log file size....476
log filenames....476
network element configuration....473
network element information
verifying....481
node information
verifying....478
origin attribute configuration....472
peer configuration....473
peer information
verifying....480
route information
verifying....478
tracing operations....475
transport configuration....474
troubleshooting configuration....482
troubleshooting connectivity....482
Diameter Base Protocol
message severity levels for tracing operations....478
Diameter base protocol statements
address
peer....1434
transport....1435
connect-actively....1520
destination....1533
diameter....1562
forwarding....1647
function
network element....1653
route....1654
host....1685
logical-system....1760
transport....1761
metric....1788
network-element....1813
origin....1837
peer
network element....1862
remote....1862
port....1874
priority....1895
realm....1927
route....1952
routing-instance
peer....1955
transport....1956

Copyright © 2013, Juniper Networks, Inc.

traceoptions....2039
transport
local....2058
peer....2058
diameter statement
Diameter base protocol....1562
diameter-instance statement
Gx-Plus....1563
JSRC....1563
PTSP....1564
directed logout
AAA....214, 288
disable statement....1564
dynamic IGMP....1565
dynamic MLD....1565
disable-calling-number-avp statement....1566
disable-failover-protocol statement....1566
disable-relay statement....1567
discovery table
ANCP global configuration....1335
ANCP neighbor configuration....1333, 1335
DNS (Domain Name System)
name server address
configuring....151
overview....150
preference order....150
DNS statements
domain-name-server....1574
domain-name-server-inet....1575
domain-name-server-inet6....1576
dns-server statement....1567
documentation
comments on....lxxii
domain map....180
AAA logical system/routing instance....184
access profile....182
address pool....184
configuring....181
domain name....187
dynamic profile....183, 190
L2TP tunnel profile....189
L2TP tunnel switch profile....190
target logical system/routing instance....186
verifying configuration....191
domain map statements
aaa-logical-system....1414
aaa-routing-instance....1415
access-profile....1421
address-pool....1438
delimiter....1529
domain....1568
dynamic-profile....1592
map....1773
mask....1775
metric....1788
padn....1846
parse-direction....1849
strip-domain....2007
target-logical-system....2013
target-routing-instance....2014
tunnel-profile....2064
domain mapping *See* domain map
domain maps *See* default
domain names
delimiters....187
domain map....187
parsing direction....188
stripping from username....189
domain statement
domain map....1568
domain-name statement
address-assignment pools....1569
DHCP local server....1570
DHCP relay agent....1572
static subscribers....1573
domain-name-server statement
DNS....1574
domain-name-server-inet statement
DNS....1575
domain-name-server-inet6 statement
DNS....1576
downstream-rate statement
ANCP....1577
L2TP....1577
drop statement
DHCP relay agent....1578
drop-policy statement
subscriber secure policy....1579
drop-profile statement
dynamic CoS....1580
RED....1580
drop-profile-map statement
dynamic CoS....1581
dscp statement
dynamic classifiers....1582
dynamic rewrite rules....1583

Copyright © 2013, Juniper Networks, Inc.

dscp-ipv6 statement
dynamic classifiers....1584
dynamic rewrite rules....1584
DSL Forum VSAs....108
ANCP....1327
supported RADIUS messages....119
DTCP *See* subscriber secure policy
duplicate accounting reports....28
duplicate accounting statements
access-profile-name....1424
duplication-vrf....1587
vrf-name....2118
duplicate clients
DHCP....212, 214, 286, 288
configuration guidelines....213, 287
duplicate-clients-on-interface statement
DHCP local server....1585
DHCP relay agent....1585
duplicate-protection statement
dynamic PPPoE....1586
duplication statement....1587
duplication-vrf statement
duplicate accounting....1587
dynamic client reconfiguration
DHCP local server
attempts configuration....246
authentication token configuration....248
behavior on failure configuration....247
configuration overview....245
preventing binding of nonsupporting clients....249
RADIUS-initiated disconnect configuration....247
requesting....248
dynamic CoS statements
buffer-size....1483
bytes statement....1488, 1494
cell-mode statement....1494
class-of-service....1503
classifiers....1503
delay-buffer-rate....1524
drop-profile....1580
drop-profile-map....1581
dscp
dynamic classifiers....1582
dynamic rewrite rules....1583
dscp-ipv6
dynamic classifiers....1584
dynamic rewrite rules....1584
dynamic-class-of-service-options....1588
excess-priority....1620
excess-rate....1621, 1622
excess-rate-high....1623
excess-rate-low....1624
forwarding-class....1647
frame-mode statement....1651
guaranteed-rate....1672
ieee-802.1
dynamic classifiers....1689
dynamic rewrite rules....1690
inet-precedence
dynamic classifiers....1699
dynamic rewrite rules....1699
interface....1710
interface-set....1726
interfaces....1733
loss-priority....1765
output-traffic-control-profile....1839
overhead-accounting....1841
priority....1896
protocol....1903
rewrite-rules....1949
scheduler....1968
scheduler-map....1968
scheduler-maps....1969
schedulers....1970
shaping-rate....1988
traffic-control-profiles....2055
transmit-rate....2057
unit....2081
vendor-specific-tags
access-loop-encapsulation....2105
vlan-tag
dynamic classifiers....2114
dynamic rewrite rules....2115
dynamic firewall filters
applying fast update filters....1193
attaching statically created
any interface type....1170
specific family type....1169
attaching with RADIUS....1172
basic syntax....1135, 1154
classic filters....1133
components....1133, 1147, 1165
configuration guidelines....1135, 1148, 1164
configuring fast update filters....1183
configuring interface-shared....1217
examples....1199

Copyright © 2013, Juniper Networks, Inc.

fast update filter example....1209
fast update filters....1146, 1150
fast update filters syntax....1149
ordering....1173
overview....1132
permitting expected traffic....1187
processing order....1134, 1147, 1165
types....1133
dynamic firewalls statements
adf....1440
family....1633
fast-update-filter....1638
filter....1640
firewall....1643
input....1702
interface-shared....1729
interface-specific....1738
match-order....1775
output....1838
post-service-filter....1874
precedence....1885
service....1977
service-filter....1979
service-set....1984
shared-name....1989
term....2016
dynamic home assignment
configuring Mobile IP....572
Dynamic Host Control Protocol *See* DHCP
dynamic IGMP statements
accounting....1427
disable....1565
group
with source....1664
without source....1664
group-limit....1668
group-policy....1669
igmp....1693
immediate-leave....1694
interface....1709
no-accounting
interface....1816
oif-map
interface....1821
passive
interface....1852
promiscuous-mode
interface....1902
source
interface....1993
ssm-map
interface....1999
static
interface....2002
version
interface....2107
dynamic IP demux interface statements
demux-source....1532
family....1634
dynamic IP demux statements
mac-validate....1770
dynamic MLD
overview....1373
dynamic MLD interface statements
accounting....1427
exclude....1629
group....1665
group-count....1667
group-increment....1667
group-limit....1669
group-policy....1670
immediate-leave....1695
no-accounting....1427
oif-map....1821
passive....1853
source....1993
source-count....1995
source-increment....1996
ssm-map....1999
static....2003
version....2108
dynamic MLD statements
disable....1565
interface....1711
mld....1792
dynamic PPP statements
authentication....1467
challenge-length....1495
chap....1497
on-demand-ip-address....1822
pap....1848
ppp-options....1882

Copyright © 2013, Juniper Networks, Inc.

dynamic PPPoE
example
over dynamic VLAN demux over aggregated Ethernet....876
over static VLAN demux over aggregated Ethernet....870
dynamic PPPoE statements
duplicate-protection....1586
dynamic-profile....1593
family....1635
max-sessions....1781
max-sessions-vsa-ignore....1784
pp0....1875
pppoe-options....1877
pppoe-underlying-options....1880
server....1974
short-cycle-protection....1990
underlying-interface....2075
unit....2079
unnumbered-address....2087
dynamic profile
dynamic VLAN interface sets....1594
versions....639
dynamic profiles
associating fast update filters....1193
associating service sets....1197
components....10
configuring basic....673
configuring for client access....679
configuring global static subscriber....502, 506
configuring services levels....681
configuring static subscriber group....505
DHCP attachment....236, 330
overview....207, 282
domain map....183, 190
examples....695, 696, 697
interface support....638
overview....207, 282
MLPPP....697
MLPPP attachment....375, 811
modifying....684
overview....638
PPP....363, 365
PPP attachment....370
PPPoE....697
PPPoE interfaces....363
predefined variables....642
router predefined variables....10
tiered service example....1379
user-defined variables....666
variable expressions....667
versioning....683
VLAN....714
dynamic profiles statements
access....1417
access-internal....1420
attribute....1463
default-value....1523
dial-options....1561
dynamic-profiles....1598
equals....1617
interface....1712
dynamic routing options....1714
interfaces....1734
keepalives....1747
mandatory....1772
metric....1787
mld....1792
multicast
dynamic routing options....1795
next-hop....1814
no-keepalives....1819
no-qos-adjust
dynamic routing options....1819
ppp-subscriber-services....1884
predefined-variable-defaults....1886
preference....1887
protocols....1904
qualified-next-hop....1914
radius....1917
route
access....1950
access-internal....1951
router-advertisement....1954
routing-instances....1961
routing-options....1963
tag
access routes....2011
CoS dynamic service profile....2012
uid....2073
unnumbered-address....2086
variables....2104
vendor-id....2105
vlan-id....2110
vlan-tags....2117
dynamic protocols
overview....1371

Copyright © 2013, Juniper Networks, Inc.

dynamic requests
RADIUS....86, 90
dynamic Router Advertisement protocol
overview....1375
dynamic router advertisement statements
autonomous....1473
current-hop-limit....1520
default-lifetime....1523
interface....1713
managed-configuration....1772
max-advertisement-interval....1776
min-advertisement-interval....1789
no-managed-configuration....1772
no-other-stateful-configuration....1837
on-link....1822
other-stateful-configuration....1837
preferred-lifetime....1890
prefix....1893
reachable-time....1926
retransmit-timer....1945
router-advertisement....1954
valid-lifetime....2103
dynamic service activation
during login....86
dynamic service sets
applying fast update filters....1197
overview....1151
dynamic stacked VLAN
password authentication....723
dynamic subscribers
interfaces statement....1734
max-sessions-vsa-ignore statement....1784
pppoe-underlying-options statement....1880
short-cycle-protection statement....1990
dynamic underlying interface statements
demux-source....1533
dynamic variables
configuring....675, 676
configuring unique identifiers....682
overview....640
dynamic VLAN
agent circuit identifier interface sets
associating with subscriber interfaces....1727
defining....1710, 1728
agent circuit identifier-based
adjusting CoS shaping rate and overhead accounting....742
bandwidth management overview....706
bandwidth management restrictions....709
configuration components....704
overview....702
verifying and managing....747
authentication....728
clearing ACI interface sets....748
configuring based on agent circuit identifier information....733
configuring dynamic subscriber interfaces based on ACI information....740
configuring dynamic underlying interfaces to use ACI information....737
configuring static underlying interfaces to use ACI information....739
defining ACI interface sets....735
dynamic profiles....714
any TPID configuration....716
any TPID example....751
stacked configuration....717
stacked VLAN example....752
standard TPID configuration....714
standard TPID example....751
using routing instances....718
Ethernet packet types
configuring....721
single-tag dynamic profiles....721
stacked dynamic profiles....722
general procedure....701
interfaces
configuring....720
flexible tagging example....753
single-tag example....752
single-tag VLAN dynamic profiles....720
stacked tagging example....753
stacked VLAN dynamic profiles....721
overview....701
password authentication....723
PPPoE options
configuring....1879
profile override....727
ranges....724
mixed VLAN....726
single-level VLAN....724
stacked VLAN....725
standard ethertype example....754
removing when no subscribers....731
verifying configuration....747

Copyright © 2013, Juniper Networks, Inc.

dynamic VLAN interface sets
configuring....1446, 1472, 1594
dynamic VLAN statements
agent-circuit-identifier....1446
auto-configure....1472
dynamic profile....1594
interface....1710
interface-set
associating....1727
defining....1728
pppoe-underlying-options....1879
dynamic VLANs....1471
dynamic-class-of-service-options statement
dynamic CoS....1588
dynamic-home-assignment statement
Mobile IP....1589
dynamic-profile parsing....1162
dynamic-profile statement
DHCP local server....1590
DHCP relay agent....1591
domain map....1592
dynamic PPPoE....1593
L2TP....1594
MLPPP....1595
PPP....1595
PPPoE service name tables....1596
static subscribers....1597
dynamic-profiles
interfaces statement....1734
dynamic IP demux....1734
dynamic-profiles statement....1598

E

effective shaping rate reporting
verifying configuration....1108
effective shaping rate, CoS
reporting....1108
effective-shaping-rate statement
CoS....1605
enable statement....1605
enable-service statement
Mobile IP....1606
enable-snmp-tunnel-statistics statement
L2TP....1607
encapsulation statement
dynamic profiles....1608
logical interfaces....1611
encapsulation types
ATM subscriber access....585
enforce-strict-scale-limit-license statement
subscriber management....1615
enhanced-policer statement....1615
entity-type statement
Mobile IP....1617
equals statement
DHCP relay agent....1616
dynamic profile variables....1617
Ethernet interfaces
dynamic VLANs....1471
OAM support on service
VLANs....709, 743, 1820
VLAN tagging....2116
ethernet-port-type-virtual statement....1618
examples
configuring IP fragment reassembly on
l2tp....421
exceed-action statement....1618
excess bandwidth distribution
MIC and MPC interfaces....1069
excess-burst-size statement....1619
excess-priority statement
dynamic CoS....1620
excess-rate statement
dynamic scheduling....1621
dynamic traffic shaping....1622
excess-rate-high statement
dynamic traffic shaping....1623
excess-rate-low statement
dynamic traffic shaping....1624
exclude statement....1625
dynamic MLD interface....1629
expressions....667
external-authority statement
DHCP local server pool matching....1629

F

fail filters
unicast RPF for subscriber interfaces
configuring....1180
configuring overview....1179
example....1219
failover-within-preference-avp statement....1630
failure-action statement
liveness detection....1631
family statement
address-assignment pools....1632
dynamic firewalls....1633
dynamic IP demux interface....1634

Copyright © 2013, Juniper Networks, Inc.

dynamic PPPoE....1635
dynamic profiles....1636
Fast Ethernet interfaces
dynamic VLANs....1471
VLAN tagging....2116
fast update filters....1132
actions....1150, 1187
adding a term once....1151
applying to interfaces....1193
associating to dynamic profiles....1193
basic syntax....1149
components....1147
configuration guidelines....1148
configuring....1183
configuring match order....1184
configuring terms....1185
conflict errors....1188, 1191
evaluating terms....1189
example....1209
implied wildcard....1190
match conditions....1150, 1186
implied wildcard....1150
names....1148
only-at-create....1151
overlapping terms....1188, 1191
overview....1146
processing order....1147
fast-update-filter statement
dynamic firewalls....1638
filter precedence....1165
filter statement
dynamic firewalls....1640
dynamic interface unit....1641
firewall....1639
filter-specific statement....1642
filters
parameterized....1153
verifying configuration....1179, 1194
firewall
fast update filter actions....1187
fast update filter match conditions....1186
firewall filter
DHCP on routers using the jdhcpd process....208
example
DHCP on routers using the jdhcpd process....265
firewall filters *See* dynamic firewall filters
classic filters....1133
configuring fast update filters....1183
fast update filters....1132, 1146
overview....1133
firewall hierarchical-policer....1163
firewall policer....1163
firewall statement
dynamic profiles....1643
font conventions....lxxi
forward-only statement
DHCP relay agent....1644
forward-snooped-clients statement
DHCP local server....1645
DHCP relay agent....1646
forwarding statement
Diameter base protocol....1647
forwarding-class statement
dynamic CoS....1647
subscriber secure policy....1648
fpc statement
MX Series routers....1649
fragment reassembly
configuring on l2tp....421
on l2tp....390
frame-mode statement
CoS statements....1651
dynamic CoS statements....1651
from statement
captive-portal-content-delivery....1652
subscriber secure policy....1653
function statement
Diameter base protocol
network element....1653
route....1654

G

gateway-name statement
tunnels
LAC....1655
LNS....1654
general authentication service
tracing operations....175
generic statement
Mobile IP....1655
Gigabit Ethernet interfaces
dynamic VLANs....1471
VLAN tagging....2116

global attributes
Gx-Plus, configuring....551
global statement
Gx-Plus....1656
grace-period statement....1656
graceful Routing Engine switchover
DHCP....254
graceful Routing Engine switchover (GRES)
L2TP....389
GRES *See* graceful Routing Engine switchover (GRES)
gres-route-flush-delay statement
subscriber management....1657
group
configuring static subscriber....504
group statement
DHCP local server....1658
DHCP relay agent....1661
dynamic IGMP
with source....1664
without source....1664
dynamic MLD interface....1665
static subscribers....1666
group-count statement
dynamic MLD interface....1667
group-increment statement
dynamic MLD interface....1667
group-limit statement
dynamic IGMP....1668
dynamic MLD interface....1669
group-policy statement
dynamic IGMP....1669
dynamic MLD interface....1670
guaranteed-rate statement
dynamic CoS....1672
Gx-Plus
configuration overview....549
Diameter AVPs....462
Diameter message sequences....543
Diameter messages....457
disabling PCRF policy control....552
global attributes
configuring....551
inclusion of IPv6 subscribers
configuring....551
interactions with the PCRF....543
number of outstanding requests
configuring....551
overview....541
partition
configuring....550
provisioning subscribers....552
tracing operations....175
verifying....478, 479
gx-plus statement
Gx-Plus....1673
Gx-Plus statements
destination-host....1536
destination-realm....1538
diameter-instance....1563
global....1656
gx-plus....1673
include-ipv6....1696
max-outstanding-requests....1780
partition....1850
provisioning-order....1906

H

hardware-address statement....1673
hello-interval statement
L2TP....1674
hierarchical policer....1214
configuration statement for
aggregate....1447
example....1214
hierarchical-policer statement....1675
hierarchical-scheduler....956
implicit-hierarchy....942, 944
hierarchical-scheduler statement....1677
high-utilization statement
address-assignment pools....1678
high-utilization-v6 statement
address-assignment pools....1679
holddown-interval statement
liveness detection....1680
home agent
configuration overview....570
home agent, Mobile IP *See* Mobile IP home agent
home-agent statement
Mobile IP
dynamic home assignment rule....1681
IP address rule....1682
networks....1683
home-agent-address statement
Mobile IP....1684
host statement
address-assignment pools....1685
Diameter base protocol....1685

Copyright © 2013, Juniper Networks, Inc.

HTTP redirect
configuring subscriber interfaces.....1229
remote operation flow.....1226, 1227
HTTP service
example configuring attached to a dynamic interface.....1237
example configuring attached to a static interface.....1235
example configuring within a service set.....1235
HTTP_redirect
example DA rewrite.....1240
example redundant multiservice.....1239

I

IA_NA
DHCPv6 local server.....230
icons defined, notice.....lxx
identification statement
tunnels.....1686
idle timeout
subscriber access.....153
idle-timeout
L2TP tunnels.....412
idle-timeout statement.....1687, 1688
ieee-802.1 statement
dynamic classifiers.....1689
dynamic rewrite rules.....1690
ietf-mode statement
ANCP.....1690
if-exceeding statement
hierarchical policer.....1692
single-rate two-color policer.....1691
IGMP
enabling.....1693, 1905
version.....2107
igmp statement
dynamic IGMP.....1693
IGMP statements
promiscuous-mode interface.....1902
ignore statement.....1694
immediate-leave statement
dynamic IGMP.....1694
dynamic MLD interface.....1695
immediate-update statement
accounting.....1696
implicit-hierarchy.....927, 930, 942, 944, 956
implicit-hierarchy statement.....1677
inactive VLANs
removing.....155
include-ipv6 statement
Gx-Plus.....1696
inet-precedence statement
dynamic classifiers.....1699
dynamic rewrite rules.....1699
inline service interfaces
configuring for L2TP LNS.....427
enabling for L2TP LNS.....427
inline services statements
bandwidth.....1474
inline-services.....1700
inline-services statement.....1700
inner-tag-protocol-id statement
dynanic VLAN interfaces.....1700
inner-vlan-id statement
dynanic VLAN interfaces.....1701
input statement
dynamic service sets.....1702
input-vlan-map statement
dynamic interfaces.....1703
instance-role statement.....1704
interface delete events
maintaining subscribers.....232, 313
interface groups
DHCP local server
configuration guidelines.....218, 291
options.....219
DHCP relay
configuration guidelines.....218, 291
options.....292
enabling static subscribers to log in.....508
forcing static subscribers to log out.....508
interface ranges
DHCP configuration guidelines.....218, 291
interface sets
applying CoS traffic shaping attributes.....996, 998, 1003
CoS traffic shaping predefined variables.....1005
unit.....2078
interface statement
DHCP local server.....1705
DHCP relay agent.....1707
dynamic CoS.....1710
dynamic IGMP.....1709
dynamic MLD.....1711
dynamic profiles.....1712

dynamic router advertisement....1713
dynamic VLAN
defining....1710
L2TP service interfaces....1714
multicast
dynamic routing options....1714
static subscribers
group....1715
username....1716
interface-client-limit statement
DHCP local server....1717
DHCP relay agent....1719
interface-delete statement
subscriber management....1720
interface-description-format statement....1721
interface-id statement....1722
interface-name statement
DHCP local server....1723
DHCP relay agent....1724
interface-set
dynamic
configuring....1065
interface-set statement
ANCP....1725
dynamic CoS....1726
dynamic VLAN
associating....1727
defining....1728
interface-shared statement
dynamic firewalls....1729
interface-specific statement
dynamic firewalls....1738
interface-traceoptions statement
DHCP local server....1730
DHCP relay agent....1730
interfaces
enabling static subscribers to log in....508
forcing static subscribers to log out....508
unit statement....2083
interfaces statement
ANCP....1732
dynamic CoS....1733
dynamic profiles....1734
subscriber secure policy....1731
Internet Group Management Protocol *See* IGMP
IP demultiplexing interface statements
unit....2076
IP fragment reassembly
configuring on l2tp....421
on l2tp....390
ip-address statement....1738
ip-address-first statement....1740
ip-reassembly-rules statement
service-set....1743
ISSU *See* see unified ISSU

J

JSRC
authorizing subscribers....493
configuration overview....491, 500
Diameter AVPs....462
Diameter message sequences....488
Diameter messages....457
interactions with the SAE....488
managing subscribers....485
partition
assigning....493
configuring....492
provisioning services....485
provisioning static subscribers....495
provisioning subscribers....494
service accounting....486
tracing operations....175
verifying....478, 479
jsrc statement
JSRC....1743
JSRC statements
accounting-order....1428
authorization-order....1470
destination-host....1535
destination-realm....1538
diameter-instance....1563
jsrc....1743
jsrc-partition....1744
partition....1850
provisioning-order....1906
service....1976
jsrc-partition statement
JSRC....1744
Juniper Networks VSAs....90
corresponding predefined variables....660
supported....98
juniper-dsl-attributes statement
RADIUS....1744
junos-subscriber-ip-address....1163

Copyright © 2013, Juniper Networks, Inc.

K

keepalive requests, fast
    subscriber-initiated....371
keepalive statement....1745
keepalives statement....1746
    dynamic profiles....1747
key statement
    Mobile IP....1748

L

L2TP
    GRES....389
l2tp
    configuring IP fragment reassembly....421
    IP fragment reassembly overview....390
L2TP (Layer 2 Tunneling Protocol)
    AVPs
        behavior at LTS....396
    configuration example
        LNS....440
    defining....379
    deleting destinations, session, statistics, tunnels....438
    destination lockout timeout....414
    destruct timeout....413
    event logging....434
    failover....388
    flags for tracing operations....437
    LNS
        configuration overview....417
    locking out destinations....414
    log file access for tracing operations....436
    log file size and number....435
    log filenames....435
    maximum retransmission count....411
    message severity levels for tracing operations....437
    peer resynchronization....388
    receive window size....412
    regular expressions for tracing operations....436
    terminology....382
    tracing operations....434
    tunnel idle timeout....412
    tunnel profile
        domain map configuration....189
    tunnel profile configuration....403
        LAC address....403
        LAC hostname....403
        LNS address....403
        LNS hostname....403
        logical system....403
        maximum sessions....403
        NAS port method....403
        password....403
        preference....403
        profile name....403
        routing instance....403
        tunnel assignment ID....403
        tunnel identifier....403
        tunnel medium....403
        tunnel type....403
    tunnel switch profile
        domain map configuration....190
    tunnel switching
        configuring....432
        overview....392
    verifying configuration....438
L2TP access concentrator. *See* LAC (L2TP access concentrator)
L2TP failover protocol
    preventing L2TP LAC from negotiating....409
L2TP LAC
    subscriber secure policy....1259
L2TP LNS
    subscriber secure policy....1259, 1272
L2TP service interfaces statements
    interface....1714
    pool....1871
    service-device-pools....1978
L2TP services
    forcing expiration of destination lockout timeouts....415
    subscriber, testing....439
    tunnels, testing....439
l2tp statement....1749
L2TP statements
    LAC
        address....1435, 1436
        advisory-options....1445
        assignment-id-format....1461
        destruct-timeout....1539
        disable-calling-number-avp....1566
        disable-failover-protocol....1566
        downstream-rate....1577
        enable-snmp-tunnel-statistics....1607
        failover-within-preference....1630
        gateway-name....1654, 1655

Copyright © 2013, Juniper Networks, Inc.

identification....1686
idle-timeout....1688
l2tp....1749
logical-system....1761
max-sessions....1783
medium....1785
nas-port-method....1805
preference....1888
remote-gateway....1940
retransmission-count-established....1944
retransmission-count-not-established....1945
routing-instance....1956
rx-connect-speed-when-equal....1967
secret....1971
source-gateway....1995
traceoptions....2041
tunnel....2061
tunnel-profile....2065
tx-address-change....2068
tx-connect-speed-method....2069
type....2070
upstream-rate....2089
weighted-load-balancing....2119
LNS
aaa-access-profile....1413
bandwidth....1474
chap....1497
destination....1534
destruct-timeout....1539
dial-options....1561
dynamic-profile....1594
enable-snmp-tunnel-statistics....1607
idle-timeout....1688
inline-services....1700
interface....1714
l2tp....1749
l2tp-access-profile....1751
lockout-timeout....1756
pap....1848
pool....1871
ppp-options....1883
retransmission-count-established....1944
retransmission-count-not-established....1945
service-device-pool....1978
service-device-pools....1978
service-interface....1980
tos-reflect....2028
traceoptions....2041
tunnel....2061
tunnel switching
avp....1473
bearer-type....1480
calling-number....1489
cisco-nas-port-info....1502
tunnel-profile....2064
tunnel-switch-profile....2066
l2tp-access-profile statement....1751
LAC (L2TP access concentrator)
address change, ignoring....413
configuration overview....402
disabling Calling Number AVP 22....407
disabling L2TP failover protocol....409
function....379
interoperation....398
NAS port method....398
Receive Speed, determining....408
Rx Connect Speed AVP
sending when transmit and receive speeds are equal....409
Transmit Speed, determining....408
tunnel assignment ID format, setting....410
tunnel name format, setting....410
tunnel selection failover configuration....406
tunnel selection methods....384
tunnel selection parameter configuration....406
weighted load balancing configuration....407
lawful intercept *See* subscriber secure policy
Layer 2 Tunneling Protocol. *See* L2TP (Layer 2 Tunneling Protocol)
layer2-unicast-replies statement....1752
lcp-renegotiation statement....1753
license
subscriber scaling....235
license requirements
address-assignment pools....175
subscriber access....6
subscriber secure policy....1246
link statement
address-assignment pools....1754
liveness detection
DHCP local server....203
DHCP relay....279
liveness-detection statement....1755

Copyright © 2013, Juniper Networks, Inc.

LNS (L2TP network server)
AAA local access profile configuration.....424
access profile configuration.....422
address-assignment pool, configuring.....424
configuration example.....440
configuration overview.....417
dynamic profile, configuring.....431
enabling inline services.....427
inline service interface configuration.....427
logical interface options configuration.....428
peer interface configuration.....426
service device pool, configuring.....430
service interface pool, configuring.....430
subscriber PPP attributes, configuring
per si interface.....419
with user group profile.....418
tunnel group, configuring.....429
user group profile configuration.....422
user group profile, configuring.....418
local-gateway address statement.....1754
local-server-group statement
DHCP relay agent.....1757
lockout-timeout statement
L2TP LNS.....1756
log files
access to Diameter base protocol.....477
access to Mobile IP.....576
collecting for Juniper Technical Support.....13
configuring ANCP trace.....1345
configuring Diameter base protocol trace.....475
configuring L2TP trace.....434
configuring Mobile IP trace.....573
configuring static subscribers trace.....509
filenames for Diameter base protocol.....476
filenames for L2TP.....435
filenames for Mobile IP.....575
filenames for subscriber management
database.....159
filenames for subscriber management session
database replication.....164
number of ANCP.....1346
number of Diameter base protocol.....476
number of L2TP.....435
number of Mobile IP.....575
number of static subscribers.....510
number of subscriber management
database.....159
number of subscriber management session
database replication.....162
profile properties.....537
size of ANCP.....1346
size of Diameter base protocol.....476
size of L2TP.....435
size of Mobile IP.....575
size of packet-triggered subscribers.....528
size of PTSP.....537
size of static subscribers.....510
size of subscriber management database.....159
size of subscriber management session
database replication.....162
logical interface (aggregate) policer
configuration statement for.....1759
logical interface statements
family.....1636
logical interface-policer statement.....1759
logical system support.....5
logical-bandwidth-policer statement.....1758
logical-interface-fpc-redundancy statement
aggregated Ethernet.....1758
logical-system statement
Diameter base protocol.....1760
transport.....1761
tunnels.....1761
logical-system-name statement
DHCP local server.....1762
DHCP relay agent.....1763
static subscribers.....1764
loose mode
MAC address validation
configuring.....773
overview.....761
loss-priority statement
dynamic CoS.....1765

M

MAC address validation
dynamic subscriber interfaces
configuring.....775
overview.....761
static subscriber interfaces
configuring.....774
subscriber interfaces
configuring.....773
underlying interfaces
configuring.....775

Copyright © 2013, Juniper Networks, Inc.

mac-address statement
access internal routes....1769
DHCP local server....1767
DHCP relay agent....1768
mac-validate statement
dynamic IP demux interface....1770
maintain-subscriber statement
subscriber management....1771
maintaining subscribers
interface delete events....232, 313
verifying configuration....233, 314
managed-configuration statement
dynamic router advertisement....1772
mandatory statement
dynamic profile variables....1772
manuals
comments on....lxxii
map statement
domain map....1773
mask statement
domain map....1775
match conditions....1165
fast update filters
implied wildcard....1150
match-direction statement
content-delivery-captive-portal....1774
IP reassembly rule....1774
match-order statement
dynamic firewalls....1775
max-advertisement-interval statement
dynamic router advertisement....1776
max-data-sessions-per-subscriber statement....1753, 1776
max-outstanding-requests statement
access....1779
Gx-Plus....1780
max-pending-accounting-stops statement
RADIUS accounting....1780
max-sessions statement
dynamic PPPoE....1781
PPPoE service name tables....1782
tunnels....1783
max-sessions-vsa-ignore statement
static and dynamic PPPoE....1784
max-withhold-time statement
RADIUS accounting....1785
maximum-discovery-table-entries statement
ANCP....1777
maximum-helper-restart-time statement
ANCP....1777
maximum-hierarchy-levels 2....925, 930
maximum-lease-time statement....1778
maximum-scheduler levels 2....940
maximum-sessions-per-tunnel statement....1778
maximum–hierarchy–levels....956
maximum–hierarchy–levels statement....1677
medium statement
tunnels....1785
method statement
liveness detection....1786
metric statement
Diameter base protocol....1788
domain map....1788
dynamic profiles....1787
MIC and MPC interfaces
burst-size....1070
excess bandwidth distribution....1069
min-advertisement-interval statement
dynamic router advertisement....1789
minimum-interval statement
liveness detection....1790, 1791
MLD
enabling....1792
mld statement
dynamic profiles....1792
MLPPP
dynamic PPP subscriber services....374
dynamic profile attachment....375, 811, 1595
dynamic profiles....697
enabling PPP subscriber services....375
hardware requirements for PPP subscriber services....374
PPP subscriber services overview....373
MLPPP statements
dynamic-profile....1595
Mobile IP
access type configuration....573
accounting method....572
authentication method....570
configuration overview....569
dynamic home assignment configuration....572
event log access....576
event logging....573
filtering trace operation output....576
flags for tracing operations....576
local authentication attributes....571
log file size....575

Copyright © 2013, Juniper Networks, Inc.

log filenames....575
message severity levels for tracing operations....577
registration request authentication....570
tracing operations....573
WiMAX operation....573
Mobile IP home agent
AAA....560
accounting....563
agent discovery....557
authentication....560
home address assignment....557
mobile node registration....560
overview....557
Mobile IP statements
access-type....1425
algorithm....1452
authenticate....1464
dynamic-home-assignment....1589
enable-service....1606
entity-type....1617
generic....1655
home-agent
dynamic home assignment rule....1681
IP address rule....1682
networks....1683
home-agent-address....1684
key....1748
mobile-ip....1793
nai....1798
order....1836
peer....1861
registration-lifetime....1929
replay-method....1942
revocation-required....1948
spi....1998
statistics....2005
timestamp-tolerance....2026
traceoptions....2045
virtual-network....2109
wimax....2119
mobile-ip statement
Mobile IP....1793
mode statement....1794
MPLS pseudowire
anchor logical tunnel....930
CoS
overview....925
CoS configuration
overview....930
hierarchical-scheduler....925
implicit-hierarchy....927, 928, 942, 944
maximum-scheduler levels 2....940
implicit-hierarchy....930
maximum-hierarchy-levels 2....925, 930
ps device-name....940, 942, 944
subscriber interfaces....925
three-level scheduling
configuring....942, 944
deployment scenario....929
logical interfaces over a pseudowire interface set....944
logical interfaces over a transport logical interface....942
overview....927, 928
Pseudowire Logical Interface Set....928
Transport Logical Interface....927
two-level scheduling
configuring....940
overview....925
multi-protocol stacking....801
multicast
configuration statements....1795
multicast statement
dynamic routing options....1795
multicast traffic *See* subscriber secure policy
separating from unicast....846
multicast-interception statement....1796
multiple address assignment
DHCPv6....230
multiple DHCPv6 addresses
predefined variable....231
multiplier statement
liveness detection....1797
Multiservices DPC
configuring PTSP....530

N

nai statement
Mobile IP....1798
name-server statement....1799
NAS port method
LAC....398
nas-identifier statement....1799
NAS-Port attribute extended format
configuring for ATM interfaces....78
configuring per physical interface....72

Copyright © 2013, Juniper Networks, Inc.

configuring per stacked VLAN....75
configuring per VLAN....74
nas-port-extended-format statement
access profiles....1800
interfaces....1802
NAS-Port-ID
configuring....54
NAS-Port-ID attribute
manual configuration....44
nas-port-id-delimiter statement....1803
nas-port-id-format statement....1804
nas-port-method statement....1805
nas-port-options statement
RADIUS options....1806
NAS-Port-Type attribute
configuring per physical interface....68
configuring per stacked VLAN....70
configuring per VLAN....69
manual configuration....45
nas-port-type statement....1809
RADIUS options....1807
neighbor
parameters for ANCP....1333
neighbor statement
ANCP
for all neighbors....1811
for unique access loop circuit
identifier....1810
neighbor-discovery-router-advertisement
statement
address-assignment pools....1811
netbios-node-type statement....1812
network element
configuring Diameter....473
network statement....1812
network-element statement
Diameter base protocol....1813
next-hop statement
dynamic profiles....1814
next-hop-service statement....1815
no-accounting statement
dynamic IGMP
interface....1816
dynamic MLD interface....1427
no-adaptation statement
liveness detection....1816
no-allow-snooped-clients statement
DHCP relay agent....1817
no-bind-on-request statement
DHCP relay agent....1818
no-keepalives statement
dynamic profiles....1819
no-managed-configuration statement
dynamic router advertisement....1772
no-other-stateful-configuration statement
dynamic router advertisement....1837
no-qos-adjust statement
dynamic routing options....1819
no-vlan-id-validate....1820
node adjustments
configuring for shaping rate....1097
non-link-protected aggregated Ethernet interfaces
manually rebalancing subscribers....830
periodically rebalancing subscribers....830
rebalancing subscribers....830
notice icons defined....lxx

O

OAM
Ethernet support on service
VLANs....709, 743, 1820
triggering ANCP....1342
oam-on-svlan statement....1820
OIF maps
separating multicast traffic....846
oif-map statement
dynamic IGMP
interface....1821
dynamic MLD interface....1821
on-demand-ip-address-statement
dynamic PPP....1822
on-link statement
dynamic router advertisement....1822
opaque DHCP options....80
option 50
DHCP local server....230
option 82
DHCP relay
auto logout....310
option 82 information
DHCP relay....325
option 82 prefix
DHCP relay....326
option 82 textual description
DHCP relay....327
option statement....1823

Copyright © 2013, Juniper Networks, Inc.

option-18 statement
stacked vlan ranges....1824
vlan ranges....1824
option-37 statement
stacked vlan ranges....1825
vlan ranges....1825
option-60 statement
DHCP local server....1826
DHCP relay agent....1827
option-82 statement
address-assignment pools....1828
DHCP local server authentication....1829
DHCP local server pool matching....1830
DHCP relay agent....1831
option-match statement....1832
option-number statement
DHCP relay agent....1833
options
RADIUS....44, 50
RADIUS server....41, 44, 45, 47, 54, 56
options per interface, VLAN, or S-VLAN
NAS-Port-Type....68
options per interface, VLAN, or stacked VLAN
configuration guidelines....65
NAS-Port extended format....72, 74, 75
NAS-Port-Type....69, 70
RADIUS server....63, 66
options statement
RADIUS....1834
order statement
accounting....1835
Mobile IP....1836
origin attributes
configuring Diameter endpoint....472
origin statement
Diameter base protocol....1837
other-stateful-configuration statement
dynamic router advertisement....1837
outbound packets....1165
output statement
dynamic service sets....1838
output-traffic-control-profile statement
dynamic CoS....1839
output-vlan-map statement
dynamic interfaces....1840
outstanding requests
Gx-Plus, configuring....551
overhead-accounting cell-mode-bytes....1084
overhead-accounting frame-mode-bytes....1084
overhead-accounting statement
dynamic CoS....1841
overrides statement
DHCP local server....1842
DHCP relay agent....1844

P

packet-triggered subscribers....518
flags for tracing operations....528, 538
log file size....528
log filenames for tracing operations....528
monitoring....539
profile properties....537
record type....538
tracing operations....527, 536
packet-triggered subscribers and policy control *See* PTSP
packet-triggered-subscribers statement....1845
packet-triggered-subscribers-partition statement....1846
padn statement
domain map....1846
pap statement....1847
dynamic PPP....1848
L2TP....1848
parameterized filter: guidelines....1164
parameterized filters....1132, 1153
components....1165
processing order....1165
syntax....1154
parameterized policers....1153
parentheses, in syntax descriptions....lxxii
parse-direction statement
domain map....1849
partition
Gx-Plus, configuring....550
JSRC, assigning....493
JSRC, configuring....492
partition statement
Gx-Plus....1850
JSRC....1850
PTSP....1851
passive statement
dynamic IGMP interface....1852
dynamic MLD interface....1853

Copyright © 2013, Juniper Networks, Inc.

password
configuring global static subscriber authentication....503
configuring static subscriber group authentication....506
stacked VLAN range....723
VLAN range....723
password statement
DHCP local server....1854
DHCP relay agent....1855
static subscribers....1856
passwords
DHCP users....238
PCRF
interactions with Gx-Plus....543
peak-burst-size statement....1857
peak-information-rate statement....1859
peer
configuring Diameter....473
peer statement
Diameter base protocol
network element....1862
remote peer....1862
Mobile IP....1861
per-interface tracing operations
DHCP local server....260, 338
DHCP relay....260, 338
physical interface policer
configuration statement for....1863
physical interfaces
dynamic VLANs....1471
VLAN tagging....2116
physical-interface-policer statement....1863
pic statement
M Series, MX Series, and T Series routers....1864
policer....1163
policer statement
configuring....1865
policer, hierarchical
configuration statement for....1675
aggregate....1447
example....1214
policers
parameterized....1153
policy statement
subscriber secure policy....1697, 1698, 1867
policy-options prefix-list....1163
policy-options statement....1868
pool statement
address-assignment pools....1869
DHCP local server....1870
L2TP service interfaces....1871
pool-match-order statement....1872
pop statement
dynanic VLAN interfaces....1873
port statement
Diameter base protocol....1874
RADIUS servers....1873
post-service-filter statement
dynamic service sets....1874
pp0 statement
dynamic PPPoE....1875
PPP
access and access-internal routes....687
access over ATM networks
configuration guidelines....589
configuration overview....587
encapsulation types....585
overview....581
delaying removal of access and access-internal routes....693, 1657
dynamic profile attachment....370, 1595
dynamic profile creation....365
dynamic profiles....363
fast keepalive requests
subscriber-initiated....371
removal of access and access-internal routes....688
verifying and managing access over ATM networks....590
verifying subscriber management configuration....372
PPP attributes
configuring for L2TP LNS subscribers
per interface....419
user group profile....418
ppp statement
group profile....1876
PPP statements
dynamic-profile....1595
PPP subscriber services
configuration example
static PPP over ATM....623
controlling order of authentication protocols....367, 1467
delaying removal of access and access-internal routes....1657

Copyright © 2013, Juniper Networks, Inc.

enabling on non-Ethernet interfaces....375
fast keepalive requests
subscriber-initiated....371
hardware requirements on non-Ethernet interfaces....374
MLPPP bundle interfaces....374
modifying the length of the CHAP challenge....369, 1495
on MLPPP, overview....373
removal of access and access-internal routes....688
ppp-options statement....1881
dynamic PPP....1882
L2TP....1883
ppp-subscriber-services statement
dynamic profiles....1884
PPPoE
access over ATM networks
configuration guidelines....589
configuration overview....587
encapsulation types....585
overview....581
dynamic access over ATM network....593
dynamic profiles....697
dynamic subscriber interfaces
benefits....883
clearing lockout....913
configuration examples....917, 919
configuration overview....887, 899
configuring the underlying interface....905
configuring with additional options....903
configuring with basic options....900
creating with PPPoE service name tables....910
limiting maximum sessions....908
locking out subscriber sessions....912
lockout for subscriber sessions overview....893
lockout time for subscriber sessions overview....896
maximum session limit guidelines....892
maximum session limit overview....890
PPPoE overview....883
unified ISSU....7
verifying the configuration....914
example
dynamic PPPoE over ATM....593
dynamic PPPoE over dynamic VLAN demux over aggregated Ethernet....876
dynamic PPPoE over static VLAN demux over aggregated Ethernet....870
static PPPoE over ATM....602
static PPPoE over static VLAN demux over aggregated Ethernet....865
service name tables
creating dynamic PPPoE subscriber interfaces....910
evaluation order for matching client information....922
static access over ATM network....602
verifying and managing access over ATM networks....590
PPPoE family
on underlying interface configuration....831
pppoe-options statement
dynamic PPPoE....1877
pppoe-tags statement
CoS....1442, 1459, 1878
pppoe-underlying-options
dynamic VLAN....1879
pppoe-underlying-options statement
static and dynamic PPPoE....1880
pre-ietf mode
ANCP global configuration....1336
ANCP neighbor configuration....1333, 1336
pre-ietf-mode statement
ANCP....1893
precedence
address assignment 44
precedence statement....1885
predefined variables....642
corresponding RADIUS attributes and VSAs....660
predefined-variable-defaults statement
dynamic CoS....1886
preference statement
dynamic profiles....1887
tunnels....1888
preferred-lifetime statement....1889
dynamic router advertisement....1890
prefix statement
address-assignment pools....1890
DHCP relay agent....1891
dynamic router advertisement....1893
premium statement
hierarchical policer....1894

Copyright © 2013, Juniper Networks, Inc.

priority statement
Diameter base protocol....1895
dynamic CoS....1896
process-inform statement
DHCP local server....1897
profile statement
subscriber access....1899
promiscuous-mode statement
IGMP
interface....1902
protocol statement
dynamic CoS....1903
subscriber secure policy....1903
protocols statement
dynamic profiles....1904
provisioning-order statement
subscriber services....1906
proxy-arp statement
subscriber interfaces....1906
proxy-mode statement....1907
ps0 statement
pseudowires....1908
pseudowire interface set....944
pseudowire logical interface device
configuring....935
pseudowire subscriber interface....935
Pseudowire Logical Interface Set....928
pseudowire statements
anchor-point....1456
ps0....1908
pseudowire-service....1909
pseudowire subscriber interfaces....923
Layer 2 circuit signaling....937
Layer 2 VPN signaling....938
service logical interface configuration....939
transport logical interface configuration....937
pseudowire subscriber logical interface
device configuration....935
pseudowire logical interface device....935
pseudowire-service statement
pseudowires....1909
PTSP
configuration overview....525, 529
configuring forward rules....534
configuring forwarding instance....534
configuring rules....531
configuring service sets....534
configuring services interface....530
configuring static policies....531
configuring static rule sets....533
configuring static rules....531
Diameter AVPs....462
Diameter message sequences....517
Diameter messages....457
flags for tracing operations....538
interactions with the SAE....517
log file size....537
log filenames for tracing operations....537
managing subscribers....516
monitoring....539
overview....515
profile properties....537
provisioning packet-triggered subscribers....518
provisioning services....516
record type....538
tracing operations....527, 536
verifying....478, 479
PTSP statements
destination-host....1536
destination-realm....1539
diameter-instance....1564
packet-triggered-subscribers....1845
packet-triggered-subscribers-partition....1846
partition....1851
push statement
dynamic VLAN interfaces....1909

Q

qos-adjust statement
ANCP....1910
qos-adjust-adsl statement
ANCP....1910
qos-adjust-adsl2 statement
ANCP....1911
qos-adjust-adsl2-plus statement
ANCP....1911
qos-adjust-sds1 statement
ANCP....1912
qos-adjust-vds1 statement
ANCP....1912
qos-adjust-vds2 statement
ANCP....1913
qualified-bum-pruning-mode statement....1913
qualified-next-hop statement
dynamic profiles....1914
queue adjustments
configuring for shaping rate....1080, 1098, 1099

Copyright © 2013, Juniper Networks, Inc.

R

RADIUS
- Acct-Off messages....25
- Acct-On messages....25
- changing default idle-timeout mapping to RADIUS code....148
- changing default session-timeout mapping to RADIUS code....148
- CoA....87
- CoS parameters for initial services
  - configuring an access dynamic profile....999
  - example....1040
  - overview....989
- dynamic requests....86, 90
- mapping terminate causes to RADIUS codes....121, 148
- options....44, 50

RADIUS accounting
- back-up options, configuring....38
- backing up during an outage....32
- duplicating in a nondefault LS:RI....28
- forcing contact with offline server....39
- monitoring backup....39

RADIUS accounting statements
- access-profile-name....1424
- accounting-backup-options....1428
- duplication-vrf....1587
- max-pending-accounting-stops....1780
- max-withhold-time....1785
- vrf-name....2118

RADIUS attribute 31
- configuring....56

RADIUS attribute 5
- configuring extended format per physical interface....72
- configuring extended format per stacked VLAN....75
- configuring extended format per VLAN....74
- configuring for ATM interfaces....78

RADIUS attribute 61
- configuring per physical interface....68
- configuring per VLAN....69, 70
- manual configuration....45

RADIUS attribute 87
- configuring....54
- manual configuration....44

RADIUS attributes....90
- configuring for ATM interfaces....78
- configuring per physical interface....68, 72
- configuring per stacked VLAN....70, 75
- configuring per VLAN....69, 74
- corresponding predefined variables....660
- defining L2TP tunnels....403
- ignoring and excluding....58
- supported....91

RADIUS dynamic request information
- verifying....90

RADIUS messages
- supported DSL Forum VSAs....119

RADIUS server
- Calling-Station-ID attribute....56
- configuring interaction with....23
- configuring parameters....36
- NAS-Port-ID attribute....44, 54
- NAS-Port-Type attribute....45
- options....41, 44, 45, 47, 54, 56
- options per interface, VLAN, or stacked VLAN....63, 66
  - configuration guidelines....65

RADIUS servers *See* subscriber secure policy
- configuration example....193
- specifying....37

radius statement
- dynamic profile variables....1917
- subscriber access....1915

RADIUS statements
- accounting-order....1428
- juniper-dsl-attributes....1744
- terminate-code....2018

RADIUS terminate codes....121

radius-coa statement
- CoS....1442, 1459, 1918

radius-disconnect statement
- DHCP local server....1919

radius-flow-tap service *See* subscriber secure policy

radius-flow-tap statement....1920

RADIUS-initiated disconnect....88
- messages....89

radius-options statement....1921, 1922

radius-server statement....1923

RADIUS-sourced DHCP options....80

range statement
- address-assignment pools....1924

rapid commit
- DHCP....228

Copyright © 2013, Juniper Networks, Inc.

rapid-commit statement
DHCP local server....1925
reachable-time statement
dynamic router advertisement....1926
realm statement
Diameter base protocol....1927
rebalance-periodic statement
aggregated Ethernet....1926
receive connect speed
equal to transmit connect speed
enabling transmission of AVP 38....409
setting L2TP....408
reconfigure statement
DHCP local server....1928
redundancy mechanisms for Virtual Chassis
module....1758
registration
Mobile IP mobile node....560
registration-lifetime statement
Mobile IP....1929
relay-agent-interface-id statement
DHCP local server....1933
DHCP relay agent....1930
username....1931
relay-agent-remote-id statement
DHCP local server statements....1934
DHCP relay agent....1932
relay-agent-subscriber-id statement
DHCP local server....1935
DHCP relay agent....1936
relay-option statement....1937
relay-option-82 statement
deleting....325, 1938
relay-server-group statement
DHCP relay agent....1939
remote-gateway statement
tunnels....1940
remote-id statement....1940
replace-ip-source-with statement....1941
replay-method statement
Mobile IP....1942
report-interface-descriptions statement 1943
request-rate statement
access....1943
restart time
ANCP global configuration....1337
resynchronization, peer
L2TP....388
retransmission
L2TP control messages....411
retransmission-count-established
statement....1944
retransmission-count-not-established
statement....1945
retransmit-timer statement
dynamic router advertisement....1945
retry statement....1946
revert-interval statement....1947
revocation-required statement
Mobile IP....1948
rewrite-rules statement
dynamic CoS....1949
route statement
access internal
dynamic profiles....1951
Diameter base protocol....1952
dynamic profiles....1950
route-suppression
DHCP local server....1953
DHCP relay agent....1953
router statement
address-assignment pools....1954
router-advertisement statement
dynamic profiles....1954
routing-instance statement
Diameter base protocol
transport....1956
Diameter base protocol peer
peer....1955
PPPoE service name tables....1957
RADIUS....1955
tunnels....1956
routing-instance-name statement
DHCP local server....1958
DHCP relay agent....1959
static subscribers....1960
routing-instances statement
dynamic profiles....1961
routing-options statement
dynamic profiles....1963
rpf-check statement....1964
rule statement
captive-portal-content-delivery....1965
IP reassembly....1966
rule-set statement
captive-portal-content-delivery....1967

Copyright © 2013, Juniper Networks, Inc.

Rx Connect Speed AVP
sending of
when transmit and receive speeds are equal....409
rx-connect-speed-when-equal statement....1967
rx-window-size
L2TP tunnels....412

S

SAE
interactions with JSRC....488
interactions with PTSP....517
service accounting with JSRC....486
scheduler statement
dynamic CoS....1968
scheduler-map statement
dynamic CoS
association with traffic-control profile....1968
scheduler-maps statement
dynamic CoS
scheduler map configuration....1969
schedulers statement
dynamic CoS....1970
secret statement
access....1971
tunnels....1971
selective traffic processing
DHCP relay....321, 323, 349, 353
send-acct-status-on-config-change statement
accounting....1972
send-release-on-delete statement
DHCP relay agent....1973
server groups
DHCP relay....328
server statement
dynamic PPPoE....1974
server-group statement
DHCP relay agent....1975
server-identifier statement
address-assignment pools....1976
service accounting
JSRC....486
service activations....1164
service logical interface
pseudowire subscriber interfaces....939
service name tables
PPPoE
evaluation order for matching client information....922
service name tables, PPPoE
creating dynamic PPPoE subscriber interfaces....910
service provisioning
packet-triggered subscribers with PTSP....518
static subscribers with JSRC....495
with JSRC....485
with PTSP....516
service sets
applying to interfaces....1197
associating to dynamic profiles....1197
dynamic....1151
service statement
dynamic service sets....1977
service accounting....1976
service VLANs
Ethernet OAM support....709, 743, 1820
service-device-pool statement
L2TP....1978
service-device-pools statement
L2TP service interfaces....1978
service-filter statement
dynamic service sets....1979
service-interface statement....1980
service-name-table statement
PPPoE underlying interface....1981
service-profile statement
DHCP local server....1982
DHCP relay agent....1983
service-set statement
dynamic service sets....1984
service-sets
verifying configuration....1198
services statement
captive-portal-content-delivery....1985
session options
subscriber access....153
session options statements
client-idle-timeout....1511
client-session-timeout....1512
session startup
authentication and accounting information....31
session timeout
subscriber access....153

Copyright © 2013, Juniper Networks, Inc.

session-mode statement
liveness detection....1986
session-options statement
access profile....1987
shaping-rate adjustments
ANCP....1097
applications....1080
for subscriber local loops....1314
configuration guidelines....1083
disabling....1104
enabling....1099
example....1114
overview....1082
multicast....1097, 1098, 1099
queues....1098, 1099
scheduler nodes....1097
shaping-rate statement
dynamic CoS....1988
shared-name statement....1989
shared-secret statement....1989
short-cycle-protection statement
static and dynamic PPPoE....1990
signaling
pseudowire subscriber interfaces....937, 938
silent failover
L2TP....388
single-rate statement....1991
sip-server-address statement....1992
sip-server-domain-name statement....1992
SNMPv3 traps
subscriber secure policy....1305
subscriber secure policy configuration....1307
source statement
dynamic IGMP
interface....1993
dynamic MLD interface....1993
source-address statement
RADIUS....1994
subscriber secure policy....1994
source-count statement
dynamic MLD interface....1995
source-gateway statement
tunnels....1995
source-increment statement
dynamic MLD interface....1996
source-ipv4-address statement....1996
source-port statement
subscriber secure policy....1997
spi statement
Mobile IP....1998
SRC
packet-triggered subscriber management with PTSP....518
SAE interactions with JSRC....488
SAE interactions with PTSP....517
static subscriber management with JSRC....495
subscriber management with JSRC....485
subscriber management with PTSP....516
ssm-map statement
dynamic IGMP
interface....1999
dynamic MLD interface....1999
stacked VLANs
configuration guidelines for RADIUS options....65
configuring RADIUS options for....66, 70
overview of configuring RADIUS options for....63
stacked-vlan-ranges statement....2000
starts-with statement
DHCP relay agent....2001
stateless firewall filters
examples
configuring enhanced mode....1176
static PPPoE
example
over static VLAN demux over aggregated Ethernet....865
static PPPoE statements
max-sessions-vsa-ignore....1784
pppoe-underlying-options....1880
short-cycle-protection....1990
static statement
dynamic IGMP
interface....2002
dynamic MLD interface....2003
static subscriber statements
access-profile....1423
aggregate-clients....1451
authentication....1468
domain-name....1573
dynamic-profile....1597
group....1666
interface
group....1715
username....1716

Copyright © 2013, Juniper Networks, Inc.

logical-system-name....1764
password....1856
routing-instance....1960
static-subscribers....2004
traceoptions....2050
user-prefix....2097
username-include....2101
static subscribers....495
configuring interface groups....504
event logging....509
flags for tracing operations....512
forcing logout....508
forcing logout for interface groups....508
global access profile....501
global authentication password....503
global dynamic profile....502, 506
global username....503
group....504
group access profile....505
group authentication password....506
group dynamic profile....505
group username....507
interfaces statement....1734
log file access for tracing operations....511
log file size and number....510
log filenames for tracing operations....510
max-sessions-vsa-ignore statement....1784
message severity levels for tracing operations....512
pppoe-underlying-options statement....1880
regular expressions for tracing operations....511
resetting login state for an interface....508
resetting login state for interface groups....508
session status verifying....509
short-cycle-protection statement....1990
tracing operations....509
static-subscribers statement....2004
statistics statement access....2005
strict mode MAC address validation configuring....773
strict mode MAC address validation overview....761
strict statement DHCP local server....2006
strip-domain statement domain map....2007
subscriber AAA information verifying....164
subscriber access configuration overview....10
subscriber access environment....4
subscriber access license requirements....6
subscriber access managing access and services....9
subscriber access operation flow....8
subscriber access overview....3
subscriber access support....5
subscriber interface statements access-concentrator....1418
subscriber interface statements address....1434
subscriber interface statements chap....1497
subscriber interface statements demux-options....1531
subscriber interface statements demux-source....1532, 1533
subscriber interface statements demux0....1530
subscriber interface statements duplicate-protection....1586
subscriber interface statements dynamic PPPoE....2079
subscriber interface statements dynamic-profile....1593
subscriber interface statements family....1634, 1635, 1636
subscriber interface statements interfaces....1734
subscriber interface statements mac-validate....1770
subscriber interface statements max-sessions....1781
subscriber interface statements max-sessions-vsa-ignore....1784
subscriber interface statements mode....1794
subscriber interface statements pap....1848
subscriber interface statements pp0....1875
subscriber interface statements ppp-options....1882
subscriber interface statements pppoe-options....1877
subscriber interface statements pppoe-underlying-options....1880
subscriber interface statements proxy-arp....1906
subscriber interface statements rpf-check....1964
subscriber interface statements server....1974
subscriber interface statements short-cycle-protection....1990
subscriber interface statements underlying-interface....2074, 2075
subscriber interface statements unit....2076, 2083
subscriber interface statements unnumbered-address....2086, 2087
subscriber interface statements vlan-tagging....2116
subscriber interfaces applying CoS traffic shaping attributes....996, 998, 1003
subscriber interfaces captive portal content delivery configuring....1229
subscriber interfaces configuring in dynamic profiles....765
subscriber interfaces CoS traffic shaping predefined variables....1005

Copyright © 2013, Juniper Networks, Inc.

example
dynamic PPPoE over dynamic VLAN demux over aggregated Ethernet.....876
dynamic PPPoE over static VLAN demux over aggregated Ethernet.....870
gigabit Ethernet VLAN.....780
gigabit Ethernet VLAN with multiple logical units.....779
gigabit Ethernet VLAN with no autonegotiation.....780
IP demux over aggregated Ethernet.....838
loopback.....780
static PPPoE over static VLAN demux over aggregated Ethernet.....865
VLAN over aggregated Ethernet.....835, 840
IP demux
configuring.....769, 770
guidelines.....760
overview.....759
IP demux over aggregated Ethernet
example.....838
overview.....816
IP demux over VLAN demux
overview.....760, 795
overview.....5, 757
PPPoE
benefits.....883
configuration examples.....917, 919
configuration overview.....887
configuring the underlying interface.....905
configuring with additional options.....903
configuring with basic options.....900
dynamic configuration overview.....899
limiting maximum sessions.....908
locking out subscriber sessions.....912
lockout for subscriber sessions overview.....893
lockout time for subscriber sessions overview.....896
maximum session limit guidelines.....892
maximum session limit overview.....890
PPPoE overview.....883
unified ISSU.....7
verifying the configuration.....913, 914
PPPoE family on underlying interface.....831
VLAN
configuring.....766
overview.....758
VLAN demux
guidelines.....760
overview.....759
VLAN demux over aggregated Ethernet
overview.....760, 795, 816
VLAN over aggregated Ethernet
configuring.....824
example.....835, 840
overview.....815
subscriber local loops
CoS shaping-rate adjustments
configuration guidelines.....1083
disabling.....1104
enabling.....1099
example.....1114
overview.....1082
CoS shaping-rate adjustments with ANCP.....1314
subscriber management
packet-triggered.....518
static.....495
with JSRC.....485
with PTSP.....516
subscriber management database
flags for tracing operations.....161
log file access for tracing operations.....160
log file size and number.....159
log filenames.....159
regular expressions for tracing operations.....160
tracing operations.....158
subscriber management session database replication
flags for tracing operations.....164
log file access for tracing operations.....163
log file size and number.....162
log filenames.....164
regular expressions for tracing operations.....163
tracing operations.....161
subscriber management statements
enforce-strict-scale-limit-license.....1615
gres-route-flush-delay.....1657
interface-delete.....1720
maintain-subscriber.....1771
subscriber-management.....2009
traceoptions.....2052
subscriber provisioning
Gx-Plus.....552
JSRC.....494

Copyright © 2013, Juniper Networks, Inc.

subscriber secure policy
configuring DTCP-initiated....1281
configuring RADIUS-initiated....1261
configuring SNMPv3 traps....1307
DTCP....1273
architecture....1267
traffic mirroring interfaces....1269
DTCP configuration....1287, 1288
L2TP LAC subscribers....1259
L2TP LNS subscribers....1259, 1272
LAES compliance....1305
license requirements....1246
multicast traffic....1253, 1272
multicast traffic configuration....1266, 1289
overview....1245
RADIUS
architecture....1248
traffic mirroring interfaces....1250
RADIUS process
logged-in subscribers....1251, 1252
RADIUS server configuration....1265
radius-flow-tap service....1262, 1282
radius-flow-tap service
configuration....1264, 1284
RADIUS-initiated....1254
SNMPv3 trap example....1307
SNMPv3 traps....1305
system resources....1281
terminating
RADIUS-initiated....1266
tunnel configuration....1263, 1283
subscriber service
activating with CLI....1121
deactivating with CLI....1121
default....1127
subscriber service session
accounting statistics....35, 1175
subscriber services
activating with CLI....1122
deactivating with CLI....1122
local activation and deactivation....1122
removing PCRF control....552
subscriber session
accounting statistics....29
subscriber session database statements
database-replication....1521
traceoptions....2054
subscriber session options
configuration overview....154
removing inactive VLANs....155
subscriber-identification statement
PTSP....2008
subscriber-management statement
subscriber management....2009
subscriber-profile statement....2010
subscribers
identifying ANCP....1334
verifying configuration....776
support, technical *See* technical support
swap statement
dynamic VLAN interfaces....2010
syntax
parameterized filters....1154
syntax conventions....lxxi

T

t1-percentage statement....2071, 2072
t2-percentage statement....2072
tag statement
access....2011
dynamic profile variables....2012
dynamic profiles access route....2011
tag-protocol-id statement
dynamic VLAN map....2012
target logical system/routing instance
domain map....186
target-logical-system statement
domain map....2013
target-routing-instance statement
domain map....2014
targeted traffic
separating multicast....846
targeted traffic distribution
Virtual Chassis....2015
targeted-distribution statement
aggregated Ethernet....2015
technical support
collecting logs for....13
contacting JTAC....lxxii
term statement
captive-portal-content-delivery....2017
fast update filters....2016
terminate-code statement....2018
tftp-server statement....2019
then statement
captive-portal-content-delivery....2020

Copyright © 2013, Juniper Networks, Inc.

three-color policer....1163
three-color-policer statement....2021
time-based accounting
Mobile IP....572
timeout statement
access....2025
DHCP local server....2024
timeouts
idle and session....153
timestamp-tolerance statement
Mobile IP....2026
token statement
DHCP local server....2027
tos-reflect statement
L2TP....2028
trace operations
collecting logs for Juniper technical support....13
filtering output for Diameter base protocol....477
filtering output for Mobile IP....576
trace statement
DHCP local server....2029
DHCP relay agent....2030
traceoptions statement
address-assignment pools....2031
ANCP....2033
captive-portal-content-delivery....2035
DHCP local server....2037
DHCP relay agent....2037
Diameter base protocol....2039
general authentication service....2031
L2TP....2041
Mobile IP....2045
PTSP....2048
static subscribers....2050
subscriber management....2052
subscriber session database replication....2054
tracing operations
address-assignment pools....175
ANCP....1345
DHCP local server....255, 333
DHCP local server interface-specific....1730
DHCP relay....255, 333
DHCP relay interface-specific....1730
Diameter base protocol....475
general authentication service....175
Gx-Plus....175
JSRC....175
L2TP....434
Mobile IP....573
packet-triggered subscribers....527, 536
PTSP....527, 536
static subscribers....509
subscriber management database....158
subscriber management session database replication....161
traffic distribution mechanisms for Virtual Chassis
targeted distribution....2015
traffic mirroring *See* subscriber secure policy
DTCP
subscriber secure policy....1267
traffic mirroring interfaces....1269
RADIUS
architecture....1248
traffic mirroring interfaces....1250
RADIUS process
at subscriber login....1251
logged-in subscribers....1252
traffic shaping
adjustment of ANCP....1337, 1339, 1341
ANCP-agent reported rates for CoS....1338
traffic-control profiles
global modification....1126
modifying....1125
per-subscriber modification....1126
traffic-control-profiles statement
dynamic CoS....2055
transmit connect speed
equal to receive connect speed
enabling transmission of AVP 38....409
setting L2TP....408
transmit-interval statement
liveness detection....2022, 2023, 2056
transmit-rate statement
dynamic CoS....2057
transport
configuring Diameter....474
transport logical interface
pseudowire subscriber interfaces....937
transport statement
Diameter base protocol
local....2058
peer....2058
trigger statement
DHCP local server....2059

Copyright © 2013, Juniper Networks, Inc.

troubleshooting subscriber access
collecting logs for Juniper Technical Support....13
trust-option-82 statement....2060
tunnel assignment ID format
L2TP LAC, setting....410
tunnel idle-timeout
L2TP....412
tunnel name format
L2TP LAC, setting....410
tunnel profile statements
nas-port-method....1805
tunnel profile, L2TP
configuration....403
domain map....189
tunnel rx-window-size
L2TP....412
tunnel selection failover
configuring for L2TP LAC....406
tunnel statement....2061
tunnels....2062
tunnel statements
address
remote gateway....1435
source gateway....1436
gateway-name
remote gateway....1654
source gateway....1655
identification....1686
logical-system....1761
max-sessions....1783
medium....1785
preference....1888
remote-gateway....1940
routing-instance....1956
secret....1971
source-gateway....1995
tunnel....2062
tunnel-profile....2065
type....2070
tunnel switching statements
avp....1473
bearer-type....1480
calling-number....1489
cisco-nas-port-info....1502
tunnel-profile....2064
tunnel-switch-profile
applying....2066
defining....2066
tunnel switching, L2TP
AVP handling....396
configuring....432
domain map configuration....190
overview....392
tunnel-group statement....2063
tunnel-profile statement
domain map....2064
L2TP tunnel switch profiles....2064
tunnels....2065
tunnel-switch-profile statement
L2TP tunnel switch profiles
applying....2066
defining....2066
two-rate statement....2067
tx-address-change statement....2068
tx-connect-speed-method statement....2069
type statement
tunnels....2070

U

uid statement
dynamic profile variables....2073
uid substitution....1158
uid-reference....1154
underlying-interface statement
ANCP interface set....2073
dynamic PPPoE....2075
dynamic profiles....2074
unicast reverse path forwarding (RPF)
dynamic profiles for subscriber interfaces
configuring....1180
fail filter for subscriber interfaces
configuring....1180
for subscriber interfaces
configuring overview....1179
example....1219
unified ISSU
DHCP access model....7
PPPoE access model....7
unified ISSU state
verifying....8
unique identifiers (UIDs)....1153, 1154
unit statement
demux interfaces....2076
dynamic CoS....2081
dynamic interface sets....2078
dynamic PPPoE....2079
interfaces....2083

Copyright © 2013, Juniper Networks, Inc.

unnumbered-address statement
dynamic PPPoE....2087
dynamic profiles....2086
untagged statement
pseudowires....2087
update-interval statement
access profile....2088
upstream-rate statement
ANCP....2089
L2TP....2089
use-interface-description statement....2090
use-option-82 statement....2091
use-primary statement
DHCP local server....2092
DHCP relay agent....2093
user group profile
configuring for L2TP LNS....418
user-defined variables....666 *See* variables, user-defined
user-group-profile statement....2094
user-prefix statement
DHCP local server....2095
DHCP relay agent....2096
static subscribers....2097
username
configuring global static subscriber....503
configuring static subscriber group....507
username information
AAA authentication....729
autosense VLANs
option 82....730, 731
username-include statement
DHCP local server....2099
DHCP relay agent....2100
stacked vlan ranges....2098
static subscribers....2101
vlan ranges....2098

V

valid-lifetime statement....2102
dynamic router advertisement....2103
variables
expressions....667
overview....640
predefined....642
user-defined....666
variables statement
dynamic profile variables....2104
variables, Junos OS predefined
dynamic CoS (schedulers)
$junos-cos-scheduler....989, 1970
$junos-cos-scheduler-bs....989
$junos-cos-scheduler-dropfile-any.989,1580
$junos-cos-scheduler-dropfile-high.989,1580
$junos-cos-scheduler-dropfile-low.989,1580
$junos-cos-scheduler-dropfile-medium-high.989,1580
$junos-cos-scheduler-dropfile-medium-low.989,1580
$junos-cos-scheduler-pri....989, 1896
$junos-cos-scheduler-tx....989, 2057
example....1040
for dynamic interface sets....996, 1003, 1005
overview....989
dynamic CoS (traffic control profiles)
$junos-cos-byte-adjust.1488,1494,1651,1841
$junos-cos-byte-adjust-cell....1494, 1841
$junos-cos-byte-adjust-frame....1651, 1841
$junos-cos-delay-buffer-rate....1524
$junos-cos-excess-priority....1620
$junos-cos-excess-rate....1622
$junos-cos-excess-rate-high....1623
$junos-cos-excess-rate-low....1624
$junos-cos-overhead-accounting....1841
$junos-cos-scheduler-excess-rate....1621
$junos-cos-shaping-mode....1841
dynamic CoS (traffic-control-profiles)
$junos-cos-guaranteed-rate....1672
$junos-cos-scheduler-map....1968
$junos-cos-scheduler-shaping-rate....1988
$junos-cos-shaping-rate....1988
example....1040
for dynamic interface sets....996, 1003, 1005
overview....989
variables, Junos predefined
corresponding RADIUS attributes and VSAs....660
defaults....1886
dynamic CoS (schedulers)
$junos-cos-scheduler-bs....1483
configuring an access dynamic profile....999
dynamic CoS (traffic-control-profiles)
configuring an access dynamic profile....999
vendor-id statement
dynamic profile variables....2105

Copyright © 2013, Juniper Networks, Inc.

vendor-specific attributes
defining L2TP tunnels....403
supported....98
vendor-specific-tags access-loop-encapsulation statement
dynamic CoS....2105
verification
aggregate route....851
version statement
dynamic IGMP
interface....2107
dynamic MLD interface....2108
liveness detection....2106
versioning
dynamic profiles....639
versions
dynamic profiles....683
Virtual Chassis
module redundancy....1758
redundancy
module....1758
targeted traffic distribution....2015
virtual-network statement
Mobile IP....2109
VLAN *See* dynamic VLAN
multi-protocol stacking....801
VLAN tagging....2116
vlan-id statement
dynamic profiles....2110
dynamic VLAN map
rewriting at ingress or egress....2111
vlan-model statement....2112
vlan-nas-port-stacked-format statement....2112
vlan-ranges statement....2113
vlan-tag statement
dynamic classifiers....2114
dynamic rewrite-rules....2115
vlan-tagging statement....2116
vlan-tags statement
dynamic profiles....2117
VLANs
configuration guidelines for RADIUS options....65
configuring RADIUS options for....66, 69
overview of configuring RADIUS options....63
removing inactive....155
VLANs, service
Ethernet OAM support....709, 743, 1820
vrf-name statement
duplicate accounting....2118
VSA 26-55
opaque DHCP options....80
RADIUS-sourced DHCP options....80
VSAs
ANCP-related DSL Forum....1327
corresponding predefined variables....660
DSL Forum....108
supported....98

W

wait-for-acct-on-ack statement
accounting....2118
walled garden
example configuring as an HTTP service rule....1234
example configuring as service filter....1233
weighted load balancing
configuring for L2TP LAC....407
weighted-load-balancing statement....2119
wimax statement
Mobile IP....2119
wins-server statement....2120
wireless roaming
Mobile IP....557

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

Index of Statements and Commands

A

aaa-access-profile statement
    L2TP LNS....1413
aaa-logical-system statement
    domain map....1414
aaa-routing-instance statement
    domain map....1415
abated-utilization statement
    address-assignment pools....1415
abated-utilization-v6 statement
    address-assignment pools....1416
access profile
    session options....153
access profile statements
    access-profile-name....1424
    accounting-backup-options....1428
    duplication-vrf....1587
    max-pending-accounting-stops....1780
    max-withhold-time....1785
    session-options....1987
    update-interval....2088
    vrf-name....2118
access statement
    dynamic profiles....1417
access-concentrator statement....1418
access-identifier statement
    ANCP....1419
access-internal statement
    dynamic profiles....1420
access-loop-id-local statement....1421
access-profile statement
    domain map....1421
    routing instances....1422
    static subscribers....1423
access-profile statements
    accounting....1426
access-profile-name statement
    duplicate accounting....1424
access-type statement
    Mobile IP....1425
accounting statement
    access profile....1426
    dynamic IGMP
        interface....1427
    dynamic MLD interface....1427
accounting-backup-options statement
    RADIUS accounting....1428
accounting-order statement
    service accounting....1428
accounting-port statement....1429
accounting-server statement....1429
accounting-session-id-format
  statement....1430, 1739
accounting-stop-on-access-deny
  statement....1430
accounting-stop-on-failure statement....1431
action statement....1432
active-server-group statement
    DHCP relay agent....1433
address statement
    Diameter base protocol
        peer....1434
        transport....1435
    interface....1434
    tunnels
        LAC....1436
        LNS....1435
address-assignment pools
    abated-utilization....1415
    abated-utilization-v6....1416
    high-utilization....1678
    high-utilization-v6....1679
    router advertisement....1811
address-assignment statement
    address-assignment pools....1437
address-change-immediate-update statement
    accounting....1438
address-pool statement
    domain map....1438
address-protection statement
    access....1439
adf statement
    dynamic firewalls....1440
adjacency-timer statement
    ANCP....1441

adjust-minimum statement....1443
adjust-percent statement....1444
adjustment-control-profiles
configuring....1106
effective shaping rate....1084
overview....1084
viewing....1107
adjustment-control-profiles statement
CoS....1442
advisory-options statement
ANCP....1445
L2TP....1445
agent-circuit-identifier statement
dynamic VLAN interface sets....1446
aggregate statement
hierarchical policer....1447
aggregate-clients statement
DHCP local server....1448
DHCP relay agent....1449
static subscribers....1451
aggregated Ethernet
targeted distribution of
subscribers....1758, 1926, 2015
algorithm statement
Mobile IP....1452
allow-snooped-clients statement
DHCP relay agent....1453
always-write-giaddr statement
DHCP relay agent....1454
always-write-option-82 statement
DHCP relay agent....1455
anchor-point statement
pseudowires....1456
ancp statement
ANCP....1457
CoS....1442, 1458, 1459
ANCP statements
access-identifier....1419
adjacency-timer....1441
advisory-options....1445
ancp....1457
downstream-rate....1577
ietf-mode....1690
interface-set....1725
interfaces....1732
juniper-dsl-attributes....1744
maximum-discovery-table-entries....1777
maximum-helper-restart-time....1777
neighbor
for all neighbors....1811
for unique access loop circuit
identifier....1810
pre-ietf-mode....1893
qos-adjust....1910
qos-adjust-adsl....1910
qos-adjust-adsl2....1911
qos-adjust-adsl2-plus....1911
qos-adjust-sds1....1912
qos-adjust-vds1....1912
qos-adjust-vds2....1913
traceoptions....2033
upstream-rate....2089
application statement....1459
CoS....1442, 1459
apply-groups statement
subscriber secure policy....1460
apply-groups-except statement
subscriber secure policy....1460
attempts statement
DHCP local server....1462
attribute statement
dynamic profile variables....1463
attributes statement....1463
authenticate statement
Mobile IP....1464
authentication statement
DHCP local server....1465
DHCP relay agent....1466
dynamic PPP....1467
static subscribers....1468
authentication-order statement
access....1469
authentication-server statement....1470
authorization-order statement
JSRC....1470
auto-configure
dynamic VLAN interface sets....1472
auto-configure statement....1471
autonomous statement
dynamic router advertisement....1473
avp statement
L2TP tunnel switch profiles....1473

B

bandwidth statement
inline services....1474

Copyright © 2013, Juniper Networks, Inc.

bandwidth-limit statement
hierarchical policer....1475
policer....1476
bandwidth-percent statement
policer....1478
bearer-type statement
L2TP tunnel switch profiles....1480
bfd statement
liveness detection method....1481
boot-file statement....1482
boot-server statement....1482
buffer-size statement
dynamic CoS....1483
burst-size-limit statement....1485
hierarchical policer....1484

C

calling-number statement
L2TP tunnel switch profiles....1489
calling-station-id-delimiter statement....1490
calling-station-id-format statement....1491
captive-portal-content-delivery statements
traceoptions....2035
challenge-length statement
dynamic PPP....1495
chap statement....1496, 1497
circuit-id statement
address-assignment pools....1498
DHCP relay agent....1499
circuit-type statement
DHCP local server....1500
DHCP relay agent....1501
cisco-nas-port-info statement
L2TP tunnel switch profiles....1502
class-of-service statement
subscriber access....1503
classifiers statement
dynamic CoS....1503
clear-on-abort statement
DHCP local server....1504
client statement....1505
client-accounting-algorithm statement
RADIUS....1506
client-authentication-algorithm statement
RADIUS....1507
client-discover-match statement
DHCP local server....1508
DHCP relay agent....1509
client-id statement....1510
DHCP relay agent....1511
client-idle-timeout statement
access profile session options....1511
client-session-timeout statement
access profile session options....1512
coa-dynamic-variable-validation....1512
coa-immediate-update statement
accounting....1513
color-aware statement....1514
color-blind statement....1515
committed-burst-size statement....1516
committed-information-rate statement....1518
concurrent-data-sessions statement....1519
connect-actively statement
Diameter base protocol....1520
CoS statements
adjustment-control-profiles....1442
ancp....1442, 1458, 1459
application....1442, 1459
effective-shaping-rate....1605
pppoe-tags....1442, 1459, 1878
radius-coa....1442, 1459, 1918
current-hop-limit statement
dynamic router advertisement....1520

D

database-replication statement
subscriber session database replication....1521
default-action statement
DHCP relay agent....1522
default-lifetime statement
dynamic router advertisement....1523
default-value statement
dynamic profile variables....1523
delay-buffer-rate statement
dynamic CoS....1524
delegated-pool statement
DHCP local server....1525
delimiter statement
DHCP local server....1526
DHCP relay agent....1528
domain map....1529
demux interfaces
unit statement....2076
demux-options statement
dynamic IP demux interface....1531

Copyright © 2013, Juniper Networks, Inc.

demux-source statement
dynamic IP demux interfaces....1532
dynamic underlying interface....1533
demux0 statement
dynamic IP demux interface....1530
demux0 statements
underlying-interface....2074
destination statement
Diameter base protocol....1533
L2TP LNS....1534
destination-address statement
cpcd....1534
subscriber secure policy....1535
destination-host statement
Gx-Plus....1536
JSRC....1535
PTSP....1536
destination-port statement
subscriber secure policy....1537, 1583
destination-prefix-list statement
captive-portal-content-delivery....1537
destination-realm statement
Gx-Plus....1538
JSRC....1538
PTSP....1539
destruct-timeout statement....1539
detection-time statement
liveness detection....1540
device-count statement
pseudowires....1541
DHCP local server statements
attempts....1462
boot-file....1482
boot-server....1482
circuit-type....1500
clear-on-abort....1504
client-discover-match....1508
client-id....1510
delegated-pool 1525
delimiter....1526
dhcp-local-server....1543
dhcpv6....1554
dns-server....1567
domain-name....1570
duplicate-clients-on-interface....1585
dynamic-profile....1590
group....1658
interface....1705
interface-client-limit....1717
interface-name....1723
ip-address-first....1740
logical-system-name....1762
mac-address....1767
option-60....1826
option-82....1829, 1830
overrides....1842
password....1854
pool....1870
pool-match-order....1872
process-inform....1897
rapid-commit....1925
relay-agent-interface-id....1933
relay-agent-remote-id....1934
relay-agent-subscriber-id....1935
route-suppression....1953
routing-instance-name....1958
service-profile....1982
timeout....2024
traceoptions....2037
trigger....2059
username-include....2099
DHCP relay agent statements
active-server-group....1433
allow-snooped-clients....1453
circuit-id....1499
circuit-type....1501
client-discover-match....1509
client-id....1511
delimiter....1528
dhcp-relay....1548
dhcpv6....1557
disable-relay....1567
domain-name....1572
duplicate-clients-on-interface....1585
dynamic-profile....1591
group....1661
interface....1707
interface-client-limit....1719
interface-name....1724
layer2-unicast-replies....1752
logical-system-name....1763
mac-address....1768
no-allow-snooped-clients....1817
option-60....1827
overrides....1844
password....1855
proxy-mode....1907

Copyright © 2013, Juniper Networks, Inc.

relay-agent-interface-id....1930
username....1931
relay-agent-remote-id....1932
relay-agent-subscriber-id....1936
relay-option....1937
relay-option-82....1938
replace-ip-source-with....1941
route-suppression....1953
routing-instance-name....1959
send-release-on-delete....1973
server-group....1975
service-profile....1983
traceoptions....2037
trust-option-82....2060
use-interface-description....2090
use-option-82....2091
use-primary....2092, 2093
user-prefix....2095, 2096
username-include....1466, 2100
dhcp-attributes statement
address-assignment pools....1542
dhcp-local-server statement....1543
dhcp-relay statement....1548
DHCPv6 local server statements
preferred-lifetime....1889
t1-percentage....2071
t2-percentage....2072
valid-lifetime....2102
dhcpv6 statement....1554, 1557
dial-options statement....1560
dynamic profiles....1561
Diameter base protocol statements
address
peer....1434
transport....1435
connect-actively....1520
destination....1533
diameter....1562
forwarding....1647
function
network element....1653
route....1654
host....1685
logical-system....1760
transport....1761
metric....1788
network-element....1813
origin....1837
peer
network element....1862
remote....1862
port....1874
priority....1895
realm....1927
route....1952
routing-instance
peer....1955
transport....1956
traceoptions....2039
transport
local....2058
peer....2058
diameter statement
Diameter base protocol....1562
diameter-instance statement
Gx-Plus....1563
JSRC....1563
PTSP....1564
disable statement....1564
dynamic IGMP....1565
dynamic MLD....1565
disable-relay statement....1567
dns-server statement....1567
domain statement
domain map....1568
domain-name statement
address-assignment pools....1569
DHCP local server....1570
DHCP relay agent....1572
static subscribers....1573
downstream-rate statement
ANCP....1577
L2TP....1577
drop statement
DHCP relay agent....1578
drop-policy statement
subscriber secure policy....1579
drop-profile statement
dynamic CoS....1580
drop-profile-map statement
dynamic CoS....1581
dscp statement
dynamic classifiers....1582
dynamic rewrite rules....1583
dscp-ipv6 statement
dynamic classifiers....1584
dynamic rewrite rules....1584

duplicate accounting statements
access-profile-name....1424
duplication-vrf....1587
vrf-name....2118
duplicate-clients-on-interface statement
DHCP local server....1585
DHCP relay agent....1585
duplicate-protection statement
dynamic PPPoE....1586
duplication statement....1587
duplication-vrf statement
duplicate accounting....1587
dynamic CoS statements
buffer-size....1483
class-of-service....1503
delay-buffer-rate....1524
drop-profile....1580
drop-profile-map....1581
dscp
dynamic classifiers....1582
dynamic rewrite rules....1583
dscp-ipv6
dynamic classifiers....1584
dynamic rewrite rules....1584
dynamic-class-of-service-options 1588
excess-priority....1620
excess-rate....1621, 1622
excess-rate-high....1623
excess-rate-low....1624
forwarding-class....1647
guaranteed-rate....1672
ieee-802.1
dynamic classifiers....1689
dynamic rewrite rules....1690
inet-precedence
dynamic classifiers....1699
dynamic rewrite rules....1699
interfaces....1733
loss-priority....1765
output-traffic-control-profile....1839
overhead-accounting....1841
priority....1896
protocol....1903
scheduler....1968
scheduler-map....1968
scheduler-maps....1969
schedulers....1970
shaping-rate....1988
traffic-control-profiles....2055
transmit-rate....2057
unit....2081
vendor-specific-tags
access-loop-encapsulation....2105
vlan-tag
dynamic classifiers....2114
dynamic rewrite rules....2115
dynamic firewalls statements
family....1633
fast-update-filter....1638
filter....1640
firewall....1643
input....1702
interface-shared....1729
interface-specific....1738
match-order....1775
output....1838
post-service-filter....1874
precedence....1885
service-filter....1979
service-set....1984
shared-name....1989
term....2016
dynamic IGMP statements
accounting....1427
disable....1565
group
with source....1664
without source....1664
group-policy....1669
igmp....1693
immediate-leave....1694
interface....1709
no-accounting
interface....1816
oif-map
interface....1821
passive
interface....1852
promiscuous-mode
interface....1902
source
interface....1993
ssm-map
interface....1999
static
interface....2002
version
interface....2107

Copyright © 2013, Juniper Networks, Inc.

dynamic IP demux interface statements
demux-source....1532
family....1634
dynamic IP demux statements
mac-validate....1770
dynamic MLD interface statements
accounting....1427
exclude....1629
group....1665
group-count....1667
group-increment....1667
group-limit....1669
group-policy....1670
immediate-leave....1695
no-accounting....1427
oif-map....1821
passive....1853
source....1993
source-count....1995
source-increment....1996
ssm-map....1999
static....2003
version....2108
dynamic MLD statements
disable....1565
interface....1711
mld....1792
dynamic PPP statements
authentication....1467
challenge-length....1495
chap....1497
on-demand-ip-address....1822
pap....1848
ppp-options....1882
dynamic PPPoE statements
duplicate-protection....1586
dynamic-profile....1593
family....1635
max-sessions....1781
max-sessions-vsa-ignore....1784
pp0....1875
pppoe-options....1877
pppoe-underlying-options....1880
server....1974
short-cycle-protection....1990
underlying-interface....2075
unit....2079
unnumbered-address....2087
dynamic profile
dynamic VLAN interface sets....1594
dynamic profiles statements
access....1417
access-internal....1420
attribute....1463
default-value....1523
dial-options....1561
dynamic-profiles....1598
equals....1617
interface....1712
dynamic routing options....1714
interfaces....1734
keepalives....1747
mandatory....1772
metric....1787
mld....1792
multicast
dynamic routing options....1795
next-hop....1814
no-keepalives....1819
no-qos-adjust
dynamic routing options....1819
ppp-subscriber-services....1884
predefined-variable-defaults....1886
preference....1887
protocols....1904
qualified-next-hop....1914
radius....1917
route
access....1950
router-advertisement....1954
routing-instances....1961
tag
access routes....2011
CoS dynamic service profile....2012
uid....2073
unnumbered-address....2086
variables....2104
vendor-id....2105
vlan-tags....2117
dynamic router advertisement statements
autonomous....1473
current-hop-limit....1520
default-lifetime....1523
interface....1713
managed-configuration....1772
max-advertisement-interval....1776
min-advertisement-interval....1789

Copyright © 2013, Juniper Networks, Inc.

no-managed-configuration....1772
no-other-stateful-configuration....1837
on-link....1822
other-stateful-configuration....1837
preferred-lifetime....1890
prefix....1893
reachable-time....1926
retransmit-timer....1945
router-advertisement....1954
valid-lifetime....2103
dynamic underlying interface statements
demux-source....1533
dynamic VLAN statements
agent-circuit-identifier....1446
auto-configure....1472
dynamic profile....1594
interface....1710
interface-set
associating....1727
defining....1728
pppoe-underlying-options....1879
dynamic-class-of-service-options statement
dynamic CoS....1588
dynamic-home-assignment statement
Mobile IP....1589
dynamic-profile statement
DHCP local server....1590
DHCP relay agent....1591
domain map....1592
dynamic PPPoE....1593
L2TP....1594
MLPPP....1595
PPP....1595
PPPoE service name tables....1596
static subscribers....1597
dynamic-profiles statement....1598

E

effective-shaping-rate statement
CoS....1605
enable statement....1605
enable-service statement
Mobile IP....1606
enable-snmp-tunnel-statistics statement
L2TP....1607
encapsulation statement
dynamic profiles....1608
logical interfaces....1611
enforce-strict-scale-limit-license statement
subscriber management....1615
enhanced-policer statement....1615
entity-type statement
Mobile IP....1617
equals statement
DHCP relay agent....1616
dynamic profile variables....1617
ethernet-port-type-virtual statement....1618
exceed-action statement....1618
excess-burst-size statement....1619
excess-priority statement
dynamic CoS....1620
excess-rate statement
dynamic scheduling....1621
dynamic traffic shaping....1622
excess-rate-high statement
dynamic traffic shaping....1623
excess-rate-low statement
dynamic traffic shaping....1624
exclude statement....1625
dynamic MLD interface....1629
external-authority statement
DHCP local server pool matching....1629

F

failure-action statement
liveness detection....1631
family statement
address-assignment pools....1632
dynamic firewalls....1633
dynamic IP demux interface....1634
dynamic profiles....1636
fast-update-filter statement
dynamic firewalls....1638
filter statement
dynamic firewalls....1640
dynamic interface unit....1641
firewall....1639
filter-specific statement....1642
firewall statement
dynamic profiles....1643
forward-only statement
DHCP relay agent....1644
forward-snooped-clients statement
DHCP local server....1645
DHCP relay agent....1646
forwarding statement
Diameter base protocol....1647

Copyright © 2013, Juniper Networks, Inc.

forwarding-class statement
dynamic CoS....1647
subscriber secure policy....1648
fpc statement
MX Series routers....1649
from statement
captive-portal-content-delivery....1652
subscriber secure policy....1653
function statement
Diameter base protocol
network element....1653
route....1654

G

gateway-name statement
tunnels
LAC....1655
LNS....1654
generic statement
Mobile IP....1655
global statement
Gx-Plus....1656
grace-period statement....1656
gres-route-flush-delay statement
subscriber management....1657
group statement
DHCP local server....1658
DHCP relay agent....1661
dynamic IGMP
with source....1664
without source....1664
dynamic MLD interface....1665
static subscribers....1666
group-count statement
dynamic MLD interface....1667
group-increment statement
dynamic MLD interface....1667
group-limit statement
dynamic IGMP....1668
dynamic MLD interface....1669
group-policy statement
dynamic IGMP....1669
dynamic MLD interface....1670
guaranteed-rate statement
dynamic CoS....1672
gx-plus statement
Gx-Plus....1673
Gx-Plus statements
destination-host....1536
destination-realm....1538
diameter-instance....1563
global....1656
gx-plus....1673
include-ipv6....1696
max-outstanding-requests....1780
partition....1850
provisioning-order....1906

H

hardware-address statement....1673
hello-interval statement
L2TP....1674
hierarchical policer....1214
hierarchical-policer statement....1675
hierarchical-scheduler....956
hierarchical-scheduler statement....1677
high-utilization statement
address-assignment pools....1678
high-utilization-v6 statement
address-assignment pools....1679
holddown-interval statement
liveness detection....1680
home-agent statement
Mobile IP
dynamic home assignment rule....1681
IP address rule....1682
networks....1683
home-agent-address statement
Mobile IP....1684
host statement
address-assignment pools....1685
Diameter base protocol....1685

I

identification statement
tunnels....1686
idle timeout
subscriber access....153
idle-timeout statement....1687, 1688
ieee-802.1 statement
dynamic classifiers....1689
dynamic rewrite rules....1690
ietf-mode statement
ANCP....1690

Copyright © 2013, Juniper Networks, Inc.

if-exceeding statement
hierarchical policer....1692
single-rate two-color policer....1691
igmp statement
dynamic IGMP....1693
IGMP statements
promiscuous-mode
interface....1902
ignore statement....1694
immediate-leave statement
dynamic IGMP....1694
dynamic MLD interface....1695
immediate-update statement
accounting....1696
include-ipv6 statement
Gx-Plus....1696
inet-precedence statement
dynamic classifiers....1699
dynamic rewrite rules....1699
inline services statements
bandwidth....1474
inline-services....1700
inline-services statement....1700
inner-tag-protocol-id statement
dynanic VLAN interfaces....1700
inner-vlan-id statement
dynanic VLAN interfaces....1701
input statement
dynamic service sets....1702
input-vlan-map statement
dynamic interfaces....1703
instance-role statement....1704
interface sets
unit....2078
interface statement
DHCP local server....1705
DHCP relay agent....1707
dynamic CoS....1710
dynamic IGMP....1709
dynamic MLD....1711
dynamic profiles....1712
dynamic router advertisement....1713
dynamic VLAN
defining....1710
L2TP service interfaces....1714
multicast
dynamic routing options....1714
static subscribers
group....1715
username....1716
interface-client-limit statement
DHCP local server....1717
DHCP relay agent....1719
interface-delete statement
subscriber management....1720
interface-description-format statement....1721
interface-id statement....1722
interface-name statement
DHCP local server....1723
DHCP relay agent....1724
interface-set statement
ANCP....1725
dynamic CoS....1726
dynamic VLAN
associating....1727
defining....1728
interface-shared statement
dynamic firewalls....1729
interface-specific statement
dynamic firewalls....1738
interface-traceoptions statement
DHCP local server....1730
DHCP relay agent....1730
interfaces
unit statement....2083
interfaces statement
ANCP....1732
dynamic CoS....1733
dynamic profiles....1734
subscriber secure policy....1731
IP demultiplexing interface statements
unit....2076
ip-address statement....1738
ip-address-first statement....1740
ip-reassembly-rules statement
service-set....1743

J

jsrc statement
JSRC....1743
JSRC statements
accounting-order....1428
authorization-order....1470
destination-host....1535

Copyright © 2013, Juniper Networks, Inc.

destination-realm....1538
diameter-instance....1563
jsrc....1743
jsrc-partition....1744
partition....1850
provisioning-order....1906
service....1976
jsrc-partition statement
JSRC....1744
juniper-dsl-attributes statement
RADIUS....1744

K

keepalive statement....1745
keepalives statement....1746
dynamic profiles....1747
key statement
Mobile IP....1748

L

L2TP service interfaces statements
interface....1714
pool....1871
service-device-pools....1978
l2tp statement....1749
L2TP statements
LAC
advisory-options....1445
downstream-rate....1577
enable-snmp-tunnel-statistics....1607
upstream-rate....2089
LNS
aaa-access-profile....1413
bandwidth....1474
chap....1497
destination....1534
dial-options....1561
dynamic-profile....1594
enable-snmp-tunnel-statistics....1607
inline-services....1700
interface....1714
lockout-timeout....1756
pap....1848
pool....1871
ppp-options....1883
service-device-pool....1978
tos-reflect....2028
l2tp-access-profile statement....1751
layer2-unicast-replies statement....1752
lcp-renegotiation statement....1753
link statement
address-assignment pools....1754
liveness-detection statement....1755
local-gateway address statement....1754
local-server-group statement
DHCP relay agent....1757
lockout-timeout statement
L2TP LNS....1756
logical interface statements
family....1636
logical interface-policer statement....1759
logical-bandwidth-policer statement....1758
logical-interface-fpc-redundancy statement
aggregated Ethernet....1758
logical-system statement
Diameter base protocol....1760
transport....1761
tunnels....1761
logical-system-name statement
DHCP local server....1762
DHCP relay agent....1763
static subscribers....1764
loss-priority statement
dynamic CoS....1765

M

mac-address statement
access internal routes....1769
DHCP local server....1767
DHCP relay agent....1768
mac-validate statement
dynamic IP demux interface....1770
maintain-subscriber statement
subscriber management....1771
managed-configuration statement
dynamic router advertisement....1772
mandatory statement
dynamic profile variables....1772
map statement
domain map....1773
mask statement
domain map....1775
match-direction statement
content-delivery-captive-portal....1774
IP reassembly rule....1774
match-order statement
dynamic firewalls....1775

Copyright © 2013, Juniper Networks, Inc.

max-advertisement-interval statement
dynamic router advertisement....1776
max-data-sessions-per-subscriber statement....1753, 1776
max-outstanding-requests statement
access....1779
Gx-Plus....1780
max-pending-accounting-stops statement
RADIUS accounting....1780
max-sessions statement
dynamic PPPoE....1781
PPPoE service name tables....1782
tunnels....1783
max-sessions-vsa-ignore statement
static and dynamic PPPoE....1784
max-withhold-time statement
RADIUS accounting....1785
maximum-discovery-table-entries statement
ANCP....1777
maximum-helper-restart-time statement
ANCP....1777
maximum-lease-time statement....1778
maximum-sessions-per-tunnel statement....1778
medium statement
tunnels....1785
method statement
liveness detection....1786
metric statement
Diameter base protocol....1788
domain map....1788
dynamic profiles....1787
min-advertisement-interval statement
dynamic router advertisement....1789
minimum-interval statement
liveness detection....1790, 1791
mld statement
dynamic profiles....1792
MLPPP statements
dynamic-profile....1595
Mobile IP statements
access-type....1425
algorithm....1452
authenticate....1464
dynamic-home-assignment....1589
enable-service....1606
entity-type....1617
generic....1655
home-agent
dynamic home assignment rule....1681
IP address rule....1682
networks....1683
home-agent-address....1684
key....1748
mobile-ip....1793
nai....1798
order....1836
peer....1861
registration-lifetime....1929
replay-method....1942
revocation-required....1948
spi....1998
statistics....2005
timestamp-tolerance....2026
traceoptions....2045
virtual-network....2109
wimax....2119
mobile-ip statement
Mobile IP....1793
mode statement....1794
multicast statement
dynamic routing options....1795
multicast-interception statement....1796
multiplier statement
liveness detection....1797

N

nai statement
Mobile IP....1798
name-server statement....1799
nas-identifier statement....1799
nas-port-extended-format statement
access profiles....1800
interfaces....1802
nas-port-id-delimiter statement....1803
nas-port-id-format statement....1804
nas-port-method statement....1805
nas-port-options statement
RADIUS options....1806
nas-port-type statement....1809
RADIUS options....1807
neighbor statement
ANCP
for all neighbors....1811
for unique access loop circuit identifier....1810

Copyright © 2013, Juniper Networks, Inc.

neighbor-discovery-router-advertisement statement
address-assignment pools....1811
netbios-node-type statement....1812
network statement....1812
network-element statement
Diameter base protocol....1813
next-hop statement
dynamic profiles....1814
next-hop-service statement....1815
no-accounting statement
dynamic IGMP interface....1816
dynamic MLD interface....1427
no-adaptation statement
liveness detection....1816
no-allow-snooped-clients statement
DHCP relay agent....1817
no-bind-on-request statement
DHCP relay agent....1818
no-keepalives statement
dynamic profiles....1819
no-managed-configuration statement
dynamic router advertisement....1772
no-other-stateful-configuration statement
dynamic router advertisement....1837
no-qos-adjust statement
dynamic routing options....1819
no-vlan-id-validate....1820

O

oam-on-svlan statement....1820
oif-map statement
dynamic IGMP interface....1821
dynamic MLD interface....1821
on-demand-ip-address-statement
dynamic PPP....1822
on-link statement
dynamic router advertisement....1822
option statement....1823
option-18 statement
stacked vlan ranges....1824
vlan ranges....1824
option-37 statement
stacked vlan ranges....1825
vlan ranges....1825
option-60 statement
DHCP local server....1826
DHCP relay agent....1827
option-82 statement
address-assignment pools....1828
DHCP local server authentication....1829
DHCP local server pool matching....1830
DHCP relay agent....1831
option-match statement....1832
option-number statement
DHCP relay agent....1833
options statement
RADIUS....1834
order statement
accounting....1835
Mobile IP....1836
origin statement
Diameter base protocol....1837
other-stateful-configuration statement
dynamic router advertisement....1837
output statement
dynamic service sets....1838
output-traffic-control-profile statement
dynamic CoS....1839
output-vlan-map statement
dynamic interfaces....1840
overhead-accounting cell-mode-bytes....1084
overhead-accounting frame-mode-bytes....1084
overhead-accounting statement
dynamic CoS....1841
overrides statement
DHCP local server....1842
DHCP relay agent....1844

P

packet-triggered-subscribers statement....1845
packet-triggered-subscribers-partition statement....1846
padn statement
domain map....1846
pap statement....1847
L2TP....1848
parse-direction statement
domain map....1849
partition statement
Gx-Plus....1850
JSRC....1850
PTSP....1851

passive statement
dynamic IGMP
interface....1852
dynamic MLD interface....1853
password statement
DHCP local server....1854
DHCP relay agent....1855
static subscribers....1856
peak-burst-size statement....1857
peak-information-rate statement....1859
peer statement
Diameter base protocol
network element....1862
remote peer....1862
Mobile IP....1861
physical-interface-policer statement....1863
pic statement
M Series, MX Series, and T Series
routers....1864
policer statement
configuring....1865
policy statement
subscriber secure policy....1697, 1698, 1867
policy-options statement....1868
pool statement
address-assignment pools....1869
DHCP local server....1870
L2TP service interfaces....1871
pool-match-order statement....1872
pop statement
dynanic VLAN interfaces....1873
port statement
Diameter base protocol....1874
RADIUS servers....1873
post-service-filter statement
dynamic service sets....1874
pp0 statement
dynamic PPPoE....1875
ppp statement
group profile....1876
PPP statements
dynamic-profile....1595
ppp-options statement....1881
L2TP....1883
ppp-subscriber-services statement
dynamic profiles....1884
pppoe-options statement
dynamic PPPoE....1877
pppoe-tags statement
CoS....1442, 1459, 1878
pppoe-underlying-options
dynamic VLAN....1879
pppoe-underlying-options statement
static and dynamic PPPoE....1880
pre-ietf-mode statement
ANCP....1893
precedence statement....1885
predefined-variable-defaults statement
dynamic CoS....1886
preference statement
dynamic profiles....1887
tunnels....1888
preferred-lifetime statement....1889
dynamic router advertisement....1890
prefix statement
address-assignment pools....1890
DHCP relay agent....1891
dynamic router advertisement....1893
premium statement
hierarchical policer....1894
priority statement
Diameter base protocol....1895
dynamic CoS....1896
process-inform statement
DHCP local server....1897
profile statement
subscriber access....1899
promiscuous-mode statement
IGMP
interface....1902
protocol statement
dynamic CoS....1903
subscriber secure policy....1903
protocols statement
dynamic profiles....1904
provisioning-order statement
subscriber services....1906
proxy-arp statement
subscriber interfaces....1906
proxy-mode statement....1907
ps0 statement
pseudowires....1908
pseudowire-service statement
pseudowires....1909
PTSP statements
destination-host....1536
destination-realm....1539

Copyright © 2013, Juniper Networks, Inc.

diameter-instance....1564
packet-triggered-subscribers....1845
packet-triggered-subscribers-partition....1846
partition....1851
push statement
dynamic VLAN interfaces....1909

Q

qos-adjust statement
ANCP....1910
qos-adjust-adsl statement
ANCP....1910
qos-adjust-adsl2 statement
ANCP....1911
qos-adjust-adsl2-plus statement
ANCP....1911
qos-adjust-sds1 statement
ANCP....1912
qos-adjust-vds1 statement
ANCP....1912
qos-adjust-vds2 statement
ANCP....1913
qualified-bum-pruning-mode statement....1913
qualified-next-hop statement
dynamic profiles....1914

R

RADIUS accounting statements
access-profile-name....1424
accounting-backup-options....1428
duplication-vrf....1587
max-pending-accounting-stops....1780
max-withhold-time....1785
vrf-name....2118
radius statement
dynamic profile variables....1917
subscriber access....1915
RADIUS statements
accounting-order....1428
juniper-dsl-attributes....1744
terminate-code....2018
radius-coa statement
CoS....1442, 1459, 1918
radius-disconnect statement
DHCP local server....1919
radius-flow-tap statement....1920
radius-options statement....1921, 1922
radius-server statement....1923
range statement
address-assignment pools....1924
rapid-commit statement
DHCP local server....1925
reachable-time statement
dynamic router advertisement....1926
realm statement
Diameter base protocol....1927
rebalance-periodic statement
aggregated Ethernet....1926
reconfigure statement
DHCP local server....1928
registration-lifetime statement
Mobile IP....1929
relay-agent-interface-id statement
DHCP local server....1933
DHCP relay agent....1930
username....1931
relay-agent-remote-id statement
DHCP local server statements....1934
DHCP relay agent....1932
relay-agent-subscriber-id statement
DHCP local server....1935
DHCP relay agent....1936
relay-option statement....1937
relay-server-group statement
DHCP relay agent....1939
remote-gateway statement
tunnels....1940
remote-id statement....1940
replace-ip-source-with statement....1941
replay-method statement
Mobile IP....1942
report-interface-descriptions statement....1943
request-rate statement
access....1943
retransmission-count-established statement....1944
retransmission-count-not-established statement....1945
retransmit-timer statement
dynamic router advertisement....1945
retry statement....1946
revert-interval statement....1947
revocation-required statement
Mobile IP....1948
rewrite-rules statement
dynamic CoS....1949

Copyright © 2013, Juniper Networks, Inc.

route statement
access internal
dynamic profiles....1951
Diameter base protocol....1952
dynamic profiles....1950
route-suppression
DHCP local server....1953
DHCP relay agent....1953
router statement
address-assignment pools....1954
router-advertisement statement
dynamic profiles....1954
routing-instance statement
Diameter base protocol
transport....1956
Diameter base protocol peer
peer....1955
PPPoE service name tables....1957
RADIUS....1955
tunnels....1956
routing-instance-name statement
DHCP local server....1958
DHCP relay agent....1959
static subscribers....1960
routing-instances statement
dynamic profiles....1961
routing-options statement
dynamic profiles....1963
rpf-check statement....1964
rule statement
IP reassembly....1966
rule-set statement
captive-portal-content-delivery....1967

S

scheduler statement
dynamic CoS....1968
scheduler-map statement
dynamic CoS
association with traffic-control
profile....1968
scheduler-maps statement
dynamic CoS
scheduler map configuration....1969
schedulers statement
dynamic CoS....1970
secret statement
access....1971
tunnels....1971
send-acct-status-on-config-change statement
accounting....1972
send-release-on-delete statement
DHCP relay agent....1973
server statement
dynamic PPPoE....1974
server-group statement
DHCP relay agent....1975
server-identifier statement
address-assignment pools....1976
service statement
dynamic service sets....1977
service accounting....1976
service-device-pool statement
L2TP....1978
service-device-pools statement
L2TP service interfaces....1978
service-filter statement
dynamic service sets....1979
service-interface statement....1980
service-name-table statement
PPPoE underlying interface....1981
service-profile statement
DHCP local server....1982
DHCP relay agent....1983
service-set statement
dynamic service sets....1984
services statement
captive-portal-content-delivery....1985
session options
subscriber access....153
session options statements
client-idle-timeout....1511
client-session-timeout....1512
session timeout
subscriber access....153
session-mode statement
liveness detection....1986
session-options statement
access profile....1987
shaping-rate statement
dynamic CoS....1988
shared-name statement....1989
shared-secret statement....1989
short-cycle-protection statement
static and dynamic PPPoE....1990
single-rate statement....1991
sip-server-address statement....1992
sip-server-domain-name statement....1992

Copyright © 2013, Juniper Networks, Inc.

source statement
dynamic IGMP
interface....1993
dynamic MLD interface....1993
source-address statement
RADIUS....1994
subscriber secure policy....1994
source-count statement
dynamic MLD interface....1995
source-gateway statement
tunnels....1995
source-increment statement
dynamic MLD interface....1996
source-ipv4-address statement....1996
source-port statement
subscriber secure policy....1997
spi statement
Mobile IP....1998
ssm-map statement
dynamic IGMP
interface....1999
dynamic MLD interface....1999
stacked-vlan-ranges statement....2000
starts-with statement
DHCP relay agent....2001
static PPPoE statements
max-sessions-vsa-ignore....1784
pppoe-underlying-options....1880
short-cycle-protection....1990
static statement
dynamic IGMP
interface....2002
dynamic MLD interface....2003
static subscriber statements
access-profile....1423
aggregate-clients....1451
authentication....1468
domain-name....1573
dynamic-profile....1597
group....1666
interface
group....1715
username....1716
logical-system-name....1764
password....1856
routing-instance....1960
static-subscribers....2004
traceoptions....2050
user-prefix....2097
username-include....2101
static-subscribers statement....2004
statistics statement
access....2005
strict statement
DHCP local server....2006
strip-domain statement
domain map....2007
subscriber interface statements
access-concentrator....1418
address....1434
chap....1497
demux-options....1531
demux-source....1532, 1533
demux0....1530
duplicate-protection....1586
dynamic PPPoE....2079
dynamic-profile....1593
family....1634, 1635, 1636
interfaces....1734
mac-validate....1770
max-sessions....1781
max-sessions-vsa-ignore....1784
mode....1794
pap....1848
pp0....1875
ppp-options....1882
pppoe-options....1877
pppoe-underlying-options....1880
proxy-arp....1906
rpf-check....1964
server....1974
short-cycle-protection....1990
underlying-interface....2074, 2075
unit....2076, 2083
unnumbered-address....2086, 2087
vlan-tagging....2116
subscriber management statements
enforce-strict-scale-limit-license....1615
gres-route-flush-delay....1657
interface-delete....1720
maintain-subscriber....1771
subscriber-management....2009
subscriber-identification statement
PTSP....2008
subscriber-management statement
subscriber management....2009
subscriber-profile statement....2010

Copyright © 2013, Juniper Networks, Inc.

swap statement
dynamic VLAN interfaces....2010

T

t1-percentage statement....2071, 2072
t2-percentage statement....2072
tag statement
dynamic profile variables....2012
dynamic profiles access route....2011
tag-protocol-id statement
dynamic VLAN map....2012
target-logical-system statement
domain map....2013
target-routing-instance statement
domain map....2014
targeted-distribution statement
aggregated Ethernet....2015
term statement
fast update filters....2016
terminate-code statement....2018
tftp-server statement....2019
then statement
captive-portal-content-delivery....2020
three-color-policer statement....2021
timeout statement
access....2025
DHCP local server....2024
timestamp-tolerance statement
Mobile IP....2026
token statement
DHCP local server....2027
tos-reflect statement
L2TP....2028
trace statement
DHCP local server....2029
DHCP relay agent....2030
traceoptions statement
address-assignment pools....2031
ANCP....2033
captive-portal-content-delivery....2035
DHCP local server....2037
DHCP relay agent....2037
Diameter base protocol....2039
general authentication service....2031
L2TP....2041
Mobile IP....2045
PTSP....2048
static subscribers....2050
subscriber management....2052
subscriber session database replication....2054
traffic-control-profiles statement
dynamic CoS....2055
transmit-interval statement
liveness detection....2022, 2023, 2056
transmit-rate statement
dynamic CoS....2057
transport statement
Diameter base protocol
local....2058
peer....2058
trigger statement
DHCP local server....2059
trust-option-82 statement....2060
tunnel statement....2061
tunnels....2062
tunnel-group statement....2063
tunnel-profile statement
domain map....2064
L2TP tunnel switch profiles....2064
tunnels....2065
tunnel-switch-profile statement
L2TP tunnel switch profiles
applying....2066
defining....2066
two-rate statement....2067
tx-address-change statement....2068
type statement
tunnels....2070

U

uid statement
dynamic profile variables....2073
underlying-interface statement
dynamic PPPoE....2075
dynamic profiles....2074
unit statement
demux interfaces....2076
dynamic CoS....2081
dynamic interface sets....2078
dynamic PPPoE....2079
interfaces....2083
unnumbered-address statement
dynamic PPPoE....2087
dynamic profiles....2086
untagged statement
pseudowires....2087

Copyright © 2013, Juniper Networks, Inc.

update-interval statement
access profile....2088
upstream-rate statement
ANCP....2089
L2TP....2089
use-interface-description statement....2090
use-option-82 statement....2091
use-primary statement
DHCP local server....2092
DHCP relay agent....2093
user-group-profile statement....2094
user-prefix statement
DHCP local server....2095
DHCP relay agent....2096
static subscribers....2097
username-include statement
DHCP local server....2099
DHCP relay agent....2100
stacked vlan ranges....2098
static subscribers....2101
vlan ranges....2098

V

valid-lifetime statement....2102
dynamic router advertisement....2103
variables statement
dynamic profile variables....2104
vendor-id statement
dynamic profile variables....2105
vendor-specific-tags access-loop-encapsulation statement
dynamic CoS....2105
version statement
dynamic IGMP interface....2107
dynamic MLD interface....2108
liveness detection....2106
virtual-network statement
Mobile IP....2109
vlan-id statement
dynamic profiles....2110
dynamic VLAN map rewriting at ingress or egress....2111
vlan-model statement....2112
vlan-nas-port-stacked-format statement....2112
vlan-ranges statement....2113
vlan-tag statement
dynamic classifiers....2114
dynamic rewrite-rules....2115
vlan-tagging statement....2116
vlan-tags statement
dynamic profiles....2117
vrf-name statement
duplicate accounting....2118

W

wait-for-acct-on-ack statement
accounting....2118
wimax statement
Mobile IP....2119
wins-server statement....2120

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.