

Junos[®] OS

Broadband Subscriber Services User Guide

Published
2019-12-20

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

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

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

Junos[®] OS Broadband Subscriber Services User Guide
Copyright © 2019 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About the Documentation | xxxiii

Documentation and Release Notes | xxxiii

Using the Examples in This Manual | xxxiii

Merging a Full Example | xxxiv

Merging a Snippet | xxxv

Documentation Conventions | xxxv

Documentation Feedback | xxxviii

Requesting Technical Support | xxxviii

Self-Help Online Tools and Resources | xxxix

Creating a Service Request with JTAC | xxxix

1

Subscriber Service Activation and Management

Subscriber Service Activation and Management | 3

Dynamic Service Management with RADIUS | 3

Using RADIUS Dynamic Requests for Subscriber Access Management | 3

Benefits of Radius Dynamic Requests | 4

Configuring RADIUS-Initiated Dynamic Request Support | 5

RADIUS-Initiated Change of Authorization (CoA) Overview | 6

CoA Messages | 7

Qualifications for Change of Authorization | 7

Message Exchange | 8

Bulk CoA Transactions | 8

Benefits of Radius-Initiated Change of Authorization | 9

RADIUS-Initiated Disconnect Overview | 10

Disconnect Messages | 10

Qualifications for Disconnect | 10

Message Exchange | 10

Benefits of Radius-Initiated Disconnects | 11

Usage Thresholds for Subscriber Services | 11

Subscriber Session Logins and Service Activation Failures Overview | 12

Service and Network Family Activation Process | 14

Configuring How Service Activation Failures Affect Subscriber Login | 18

Error-Cause Codes (RADIUS Attribute 101) for Dynamic Requests | 19

Verifying RADIUS Dynamic-Request Statistics | 20

Service Activation and Deactivation Using the CLI Instead of RADIUS | 20

CLI-Activated Subscriber Services | 20

Local and Remote Service Activation and Deactivation Using the CLI | 21

Management of Subscriber Services with Multiple Instances | 25

Subscriber Services with Multiple Instances Overview | 26

Subscriber Service Instances and Service Parameters | 26

CLI Deactivation of Subscriber Services with Multiple Instances | 26

Subscriber Services with Multiple Instances in RADIUS Accounting Messages | 27

Deactivating a Single Instance of a Subscriber Service | 28

Deactivating All Instances of a Subscriber Service | 31

Verifying Subscriber Services with Multiple Instances | 34

2

Configuring Dynamic Class of Service

CoS for Subscriber Access and Interfaces Overview | 39

CoS for Subscriber Access Overview | 39

Guidelines for Configuring Dynamic CoS for Subscriber Access | 40

Configuration Guidelines for Hierarchical CoS and Per-Unit Scheduling | 40

Configuration Guidelines for Dynamic Scheduling and Queuing | 41

Configuration Guidelines for Dynamic Classifiers and Rewrite Rules | 41

CoS for Aggregated Ethernet Subscriber Interfaces Overview | 44

CoS for PPPoE Subscriber Interfaces Overview | 45

Configuring Scheduling and Shaping for Subscriber Access | 47

Configuring Traffic Scheduling and Shaping for Subscriber Access | 47

Configuring Static Traffic Shaping and Scheduling Parameters in a Dynamic Profile | 48

Configuring Dynamic Traffic Shaping and Scheduling Parameters in a Dynamic Profile | 49

Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers | 50

Using the CLI to Globally Modify a Traffic-Control Profile Currently Applied to Multiple Subscribers | 51

Using the CLI to Modify a Traffic-Control Profile for a Specific Current Subscriber | 51

Configuring Schedulers in a Dynamic Profile for Subscriber Access | 52

Configuring Static Schedulers in a Dynamic Profile | 53

Configuring Dynamic Schedulers with Variables in a Dynamic Profile | 54

Configuring a Combination of Static and Dynamic Scheduler Parameters in a Scheduler Definition | 56

Configuring Scheduler and Scheduler Map Sharing | 60

Example: Providing Unique Rate Configurations for Schedulers in a Dynamic Profile | 62

Example: Configuring Aggregate Scheduling of Queues for Residential Subscribers on Static IP Demux Interfaces | 63

Verifying the Scheduling and Shaping Configuration for Subscriber Access | 66

Configuring Hierarchical CoS Scheduling on MPLS Ethernet Pseudowire Subscriber Interfaces | 67

Enhanced Subscriber Management Subscriber Logical Interfaces or Interface Sets Over Underlying Logical Interfaces for a CoS scheduler Hierarchy | 67

Enhanced Subscriber Management Subscriber Logical Interfaces or Interface Sets Over MPLS Pseudowires for a CoS scheduler Hierarchy | 70

Configuring Layer 2 Subscriber Logical Interfaces for CoS Hierarchical Schedulers Using Dynamic Profiles for Differentiating Home and Access Node Networks | 73

Example: Configuring Layer 2 Subscriber Logical Interfaces for CoS Hierarchical Schedulers Using Static CoS for Differentiating Home and Access Node Networks | 79

Allocating Dedicated Queues for Each Logical Interface Using Per-Unit Scheduling | 85

Hardware Requirements for Dynamic Per-Unit Scheduling | 85

Configuring Per-Unit Scheduling in a Dynamic Profile | 86

Example: Configuring Per-Unit Scheduling for Subscriber Access | 88

Configuring Dedicated Queue Scaling with Hierarchical CoS or Per-Unit Scheduling | 101

Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview | 101

Queue Scaling for MPCs | 101

Managing Remaining Queues | 102

Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces | 104

Configuring the Maximum Number of Queues for MIC and MPC Interfaces | 104

Configuring Remaining Common Queues on MIC and MPC Interfaces | 105

Verifying the Number of Dedicated Queues Configured on MIC and MPC Interfaces | 106

Shaping Downstream Traffic Based on Frames or Cells | 109

Bandwidth Management for Downstream Traffic in Edge Networks Overview | 109

Effective Shaping Rate | 110

Shaping Modes | 110

Byte Adjustments | 110

Relationship with Other CoS Features | 111

Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 111

Example: Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 113

Managing Traffic with Different Encapsulations | 113

Managing Downstream Cell-Based Traffic | 115

Configuring Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 117

Example: Configuring Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 118

Managing Traffic with Different Encapsulations | 118

Managing Downstream Cell-Based Traffic | 120

Setting Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags | 121

CLI Interaction with PPPoE Vendor-Specific Tags | 121

RADIUS Interaction with PPPoE Vendor-Specific Tags | 122

ANCP Interaction with PPPoE Vendor-Specific Tags | 122

Multicast QoS Adjustment Interaction with PPPoE Vendor-Specific Tags | 122

Shaping Rate Restrictions | 122

Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces | 123

Reporting the Effective Shaping Rate for Subscribers | 124

Verifying the Effective Shaping Rate Reporting Configuration | 124

Applying CoS to Households or Individual Subscribers Using ACI-Based Dynamic VLANs | 127

Applying CoS Attributes to VLANs Using Agent-Circuit-Identifiers | 127

Agent Circuit Identifier-Based Dynamic VLANs Bandwidth Management Overview | 130

CoS Shaping Rate Adjustment | 131

CoS Overhead Accounting Adjustment | 131

Dynamic Profiles and Adjustment of CoS Shaping Rate and Overhead Accounting | 132

Guidelines for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting | 133

Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ACI Interface Sets | 134

Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs | 135

Applying CoS to Households or Individual Subscribers Using Access Line Identifier Dynamic VLANs | 137

Applying CoS Attributes to VLANs Using Access-Line Identifiers | 137

Bandwidth Management Overview for Dynamic VLANs Based on Access-Line Identifiers | 140

CoS Shaping Rate Adjustment | 141

CoS Overhead Accounting Adjustment | 141

Dynamic Profiles and Adjustment of CoS Shaping Rate and Overhead Accounting | 142

Guidelines for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting | 143

Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ALI Interface Sets | 144

Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Dynamic VLANs Based on Access-Line Identifiers | 145

Managing Excess Bandwidth Distribution and Traffic Bursts | 147

Excess Bandwidth Distribution on MIC and MPC Interfaces Overview | 147

Traffic Burst Management on MIC and MPC Interfaces Overview | 148

Guidelines for Configuring the Burst Size | 148

How the System Calculates the Burst Size | 150

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces | 151

Applying CoS Using Parameters Received from RADIUS | 155

Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS | 155

Dynamic Configuration of Initial CoS in Access Profiles | 156

Predefined Variables for Dynamic Configuration of Initial Traffic Shaping | 156

Predefined Variables for Dynamic Configuration of Initial Scheduling and Queuing | 157

Changing CoS Services Overview | 160

Types of CoS Variables Used in a Service Profile | 160

Static and Dynamic CoS Configurations | 161

Scenarios for Static and Dynamic Configuration of CoS Parameters | 161

CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview | 164

Supported Network Configurations | 164

Traffic-Control Profiles in Subscriber Interface Dynamic Profiles | 164

CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets and Member Subscriber Sessions | 165

Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions | 166

Configuring Initial CoS Parameters Dynamically Obtained from RADIUS | 167

Configuring Static Default Values for Traffic Scheduling and Shaping | 168

Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions | 170

CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets | 172

Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS | 178

Modifying a Subscriber's Shaping Characteristics After a Subscriber is Instantiated | 183

CoS Adjustment Control Profiles Overview | 183

Applications and Associated Algorithms in Adjustment Control Profiles | 185

CoS Shaping Rate Fallback Behavior | 186

Configuring CoS Adjustment Control Profiles | 186

Verifying the CoS Adjustment Control Profile Configuration | 188

Applying CoS to Groups of Subscriber Interfaces | 189

CoS for Interface Sets of Subscribers Overview | 189

Guidelines for Configuring Dynamic Interface Sets in a Subscriber Access Network | 189

Configuring an Interface Set of Subscribers in a Dynamic Profile | 192

Example: Configuring a Dynamic Interface Set of VLAN Subscribers | 193

Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile | 211

Applying CoS to Subscriber Interfaces | 217

Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile | 217

Applying Minimal Shaping and Scheduling to Remaining Subscriber Traffic | 218

Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile | 219

Applying a Classifier to a Subscriber Interface in a Dynamic Profile | 221

Configuring Dynamic Filters and Policers

Dynamic Firewall Filters Overview | 225

Understanding Dynamic Firewall Filters | 226

Defining Dynamic Filter Processing Order | 227

Configuring Static Firewall Filters That Are Dynamically Applied | 229

Classic Filters Overview | 229

Classic Filter Types | 230

Classic Filter Components | 230

Classic Filter Processing | 230

Guidelines for Creating and Applying Classic Filters for Subscriber Interfaces | 231

Basic Classic Filter Syntax | 232

Examples: Configuring Static Filters | 233

Streamlining Processing of Chains of Static Filters | 237

Configuring Firewall Filter Bypass | 237

Example: Bypassing Firewall Filters | 238

Dynamically Attaching Static or Fast Update Filters to an Interface | 245

Dynamically Attaching Statically Created Filters for a Specific Interface Family Type | 245

Dynamically Attaching Statically Created Filters for Any Interface Type | 246

Configuring Filters That Are Created Dynamically | 249

Parameterized Filters Overview | 249

Unique Identifiers for Firewall Variables | 250

Configuring Unique Identifiers for Parameterized Filters | 253

Sample Dynamic-Profile Configuration for Parameterized Filters | 254

Dynamic Profile After UID Substitutions for Parameterized Filters | 256

Multiple Parameterized Filters | 258

Parameterized Filter Processing Overview | 259

Parameterized Filters Configuration Considerations | 260

- Subscriber IP Address | 261

- Interaction with Static Configuration | 261

- Interface-Specific Dynamic Service Filters | 261

- Service Session Support | 261

- Filter Naming Conventions | 261

Guidelines for Creating and Applying Parameterized Filters for Subscriber Interfaces | 262

Parameterized Filter Match Conditions for IPv4 Traffic | 263

Parameterized Filter Match Conditions for IPv6 Traffic | 270

Parameterized Filter Nonterminating and Terminating Actions and Modifiers | 277

Firewall Filter Match Conditions for Protocol-Independent Traffic in Dynamic Service Profiles | 284

Firewall Filter Terminating and Nonterminating Actions for Protocol-Independent Traffic in Dynamic Service Profiles | 285

Interface-Shared Filters Overview | 289

Dynamically Attaching Filters Using RADIUS Variables | 290

Example: Implementing a Filter for Households That Use ACI-Based VLANs | 292

Example: Dynamic-Profile Parsing | 294

Example: Firewall Dynamic Profile | 295

Example: Configuring a Filter to Exclude DHCPv6 and ICMPv6 Control Traffic for LAC Subscriber | 296

Using Ascend Data Filters to Implement Firewalls Based on RADIUS Attributes | 303

Ascend-Data-Filter Policies for Subscriber Management Overview | 303

- Filter Naming Conventions | 304

- Use of Multiple Sessions with Ascend-Data-Filters on an Interface | 304

- Optional ADF Filter Requirement for Some Subscribers | 305

- Ascend-Data-Filter Attribute Fields | 305

- Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions | 308

- Example: Configuring Dynamic Ascend-Data-Filter Support for Subscriber Access | 310

- Example: Configuring Static Ascend-Data-Filter Support for Subscriber Access | 315

- Verifying and Managing Dynamic Ascend-Data-Filter Policy Configuration | 320

Configuring Fast Update Filters to Provide More Efficient Processing Over Classic Static Filters | 323

- Fast Update Filters Overview | 323

- Fast Update Filter Components | 324

- Fast Update Filter Processing | 325

- Fast Update Filter Names | 326

- Guidelines for Creating and Applying Fast Update Filters | 326

- Basic Fast Update Filter Syntax | 327

- Configuring Fast Update Filters | 328

- Example: Configuring Fast Update Filters for Subscriber Access | 330

- Match Conditions and Actions in Fast Update Filters | 331

- Match Conditions | 331

- Actions | 332

- Adding Terms Only Once | 332

- Configuring the Match Order for Fast Update Filters | 333

- Fast Update Filter Match Conditions | 334

- Fast Update Filter Actions and Action Modifiers | 335

- Configuring Terms for Fast Update Filters | 336

- Configuring Filters to Permit Expected Traffic | 337

- Avoiding Conflicts When Terms Match | 338

- How the Router Evaluates Terms in a Filter | 339

- Using Implied Wildcards | 340

- Conflict Caused by Overlapping Ranges | 342

- Associating Fast Update Filters with Interfaces in a Dynamic Profile | 345

Defending Against DoS and DDoS Attacks Using Unicast RPF and Fail Filters | 347

Unicast RPF in Dynamic Profiles for Subscriber Interfaces | 347

Unicast RPF in Dynamic Profiles for Subscriber Interfaces | 347

Configuring Unicast RPF in Dynamic Profiles for Subscriber Interfaces | 348

Configuring a Fail Filter for Unicast RPF in Dynamic Profiles for Subscriber Interfaces | 349

Example: Configuring Unicast RPF in a Dynamic Profile on MX Series Routers | 350

Improving Scaling and Performance of Filters on Static Subscriber Interfaces | 359

Firewall Filters and Enhanced Network Services Mode Overview | 359

Configuring a Filter for Use with Enhanced Network Services Mode | 362

Configuring Dynamic Service Sets | 365

Dynamic Service Sets Overview | 365

Associating Service Sets with Interfaces in a Dynamic Profile | 366

Verifying and Managing Service Sets Information | 367

Configuring Rate-Limiting Premium and Non-Premium Traffic on an Interface Using Hierarchical Policers | 369

Methods for Regulating Traffic by Applying Hierarchical Policers | 369

Hierarchical Policer Applied as Filter Action | 372

Example: Configuring Hierarchical Policers to Limit Rates of Services in a Static Environment | 373

Monitoring and Managing Firewalls for Subscriber Access | 389

Verifying and Managing Firewall Filter Configuration | 389

Enhanced Policer Statistics Overview | 390

4

Configuring Dynamic Multicast**Configuring Dynamic IGMP to Support IP Multicasting for Subscribers | 393**

Dynamic IGMP Configuration Overview | 393

Subscriber Management IGMP Model Overview | 394

Configuring Dynamic DHCP Client Access to a Multicast Network | 395

Example: IGMP Dynamic Profile | 397

Configuring SSM Mapping for Dynamic IGMP and MLD | 399

Configuring Dynamic MLD to Enable Subscribers to Access Multicast Networks | 401

Dynamic MLD Configuration Overview | 401

5

Configuring Application-Aware Policy Control and Reporting**Configuring Application-Aware Policy Control | 405**

Understanding Application-Aware Policy Control for Subscriber Management | 406

Benefits | 406

Understanding PCC Rules for Subscriber Management | 407

Application Filters | 408

Service Data Flow Filters | 408

PCC Action Profiles | 409

Configuring Application-Aware Policy Control for Subscriber Management | 410

Installing Services Packages for Subscriber Management Application-Aware Policy Management | 411

Configuring Service Data Flow Filters | 412

Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416

Configuring Policy and Charging Control Rules | 418

Configuring a Policy and Charging Control Rulebase | 421

Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management | 423

Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control | 424

Configuring PCC Rule Activation in a Subscriber Management Dynamic Profile | 426

Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management | 428

Configuring Application Identification | 431

Application Identification Overview | 431

Downloading and Installing Predefined Junos OS Application Signature Packages | 432

Improving the Application Traffic Throughput | 434

Configuring Custom Application Signatures | 435

Uninstalling a Predefined Junos OS Application Signature Package | 440

Configuring Reporting for Application-Aware Data Sessions | 441

Logging and Reporting Function for Subscribers | 441

Log and Report Control | 442

Templates | 442

HTTP Transaction Logging | 447

Log Dictionary for Template Types | 448

Configuring Logging and Reporting for Subscriber Management | 459

Installing Services Packages for Subscriber Management Logging and Reporting | 459

Configuring an LRF Profile for Subscribers | 460

Configuring the LRF Profile Name | 461

Configuring Policy-Based Logging | 461

(Optional) Configuring HTTP Transaction Logging | 461

Configuring Collectors | 462

Configuring Templates | 463

Configuring Logging and Reporting Rules | 465

Applying Logging and Reporting Configuration to a Subscriber Management Service Set | 467

Configuring the Activation of an LRF Rule by a PCC Rule | 468

6

Configuring HTTP Redirect Services

Configuring Captive Portal Content Delivery Services for Redirected Subscribers | 473

HTTP Redirect Service Overview | 473

Services-Card-Based Captive Portal | 476

MS-MPC-Based Captive Portal | 476

MX-SPC3 Services Card-Based Captive Portal | 477

Walled Garden Configured as a Service Filter | 477

Routing Engine-Based Captive Portal | 477

Converged Service Provisioning for HTTP Redirect Services	477
Static Service Provisioning for HTTP Redirect Services	478
Remote HTTP Redirect Server Operation Flow	480
Local HTTP Redirect Server Operation Flow	482
Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services	484
Configuring a Walled Garden as a Firewall Service Filter	485
Configuring HTTP Redirect for Local and Remote Redirect Servers	488
Configuring the Service Profile and the Service Set to Associate the Service Profile with a Service Interface	490
Attaching a CPCD Service Set and Service Filter to a Logical Interface	491
Installing a Service Package for CPCD Service	492
Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services	494
Configuring a Walled Garden as a Firewall Service Filter	495
Configuring HTTP Redirect for Local and Remote Redirect Servers	498
Configuring Parameterization for the Redirect URL	500
Configuring the Service Set to Associate the Service Profile with a Service Interface	501
Attaching a CPCD Service Set and Service Filter to a Dynamic Logical Interface	502
Installing a Service Package for CPCD Service	504
Configuring Routing Engine-Based, Static HTTP Redirect Services	505
Configuring a Walled Garden as a Firewall Service Filter	507
Configuring HTTP Redirect for Local and Remote Redirect Servers	510
Configuring the Service Profile and the Service Set to Associate the Service Profile with a Service Interface	511
Attaching a CPCD Service Set and Service Filter to a Logical Interface	513
Inserting GET Header Tags That the HTTP Server Can Use to Control Content Access	514
Configuring Routing Engine-Based, Converged HTTP Redirect Services	519
Configuring a Walled Garden as a Firewall Service Filter	520
Configuring HTTP Redirect for Local and Remote Redirect Servers	523
Configuring Parameterization for the Redirect URL	525
Configuring the Service Set to Associate the Service Profile with a Service Interface	526
Attaching a CPCD Service Set and Service Filter to a Dynamic Logical Interface	528
Adding Subscriber Information to HTTP Redirect URLs	529
How to Automatically Remove the HTTP Redirect Service After the Initial Redirect	532
Example: Configuring HTTP Redirect Services Using a Next-Hop Method and Attaching It to a Static Interface	534

Configuring Subscriber Secure Policy

Configuring Subscriber Secure Policy Traffic Mirroring Overview | 553

Subscriber Secure Policy Overview | 553

Support for Intercepting Both Layer 2 and Layer 3 Datagrams | 554

Traffic Filtering for DTCP-Initiated Subscriber Secure Policy Mirrored Traffic | 554

Mirroring-Related Event Reporting | 554

Support for L2TP Subscribers | 554

Junos OS Service for Subscriber Secure Policy Traffic Mirroring | 555

Protection of SSP Data when a Core Error is Generated | 555

Subscriber Secure Policy Licensing Requirements | 556

Configuring RADIUS-Initiated Subscriber Secure Policy Traffic Mirroring | 557

RADIUS-Initiated Subscriber Secure Policy Overview | 557

Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS | 558

RADIUS-Initiated Traffic Mirroring Interfaces | 560

RADIUS-Initiated Traffic Mirroring Process at Subscriber Login | 563

RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers | 564

RADIUS Attributes Used for Subscriber Secure Policy | 566

Triggering Subscriber Secure Policy for Subscribers on Dynamic Authenticated VLANs | 567

Using the Packet Header to Track Subscribers on the Mediation Device | 567

Format of the Mirror Header Values Used to Track Subscribers and Subscriber Sessions | 570

4-Byte Format | 570

8-Byte Format | 571

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572

Guidelines for Configuring Subscriber Secure Policy Mirroring | 573

Configuring Support for Subscriber Secure Policy Mirroring | 574

Configuring RADIUS Server Support for Subscriber Secure Policy Mirroring | 577

Terminating RADIUS-Initiated Subscriber Traffic Mirroring | 578

Configuring DTCP-Initiated Subscriber Secure Policy Traffic Mirroring | 579

DTCP-Initiated Subscriber Secure Policy Overview | 579

Subscriber Secure Policy Traffic Mirroring Architecture Using DTCP | 580

DTCP-Initiated Traffic Mirroring Interfaces | 582

DTCP-Initiated Traffic Mirroring Process | 584

DTCP Messages Used for Subscriber Secure Policy	585
Packet Header for Mirrored Traffic Sent to Mediation Device	586
Format of the Mirror Header Values Used to Track Subscribers and Subscriber Sessions	588
Manually Setting the Session-ID and Intercept ID in Packet Headers	589
Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview	590
Guidelines for Configuring Subscriber Secure Policy Mirroring	591
Configuring Support for Subscriber Secure Policy Mirroring	592
Configuring the Mediation Device as a User on the Router	595
Configuring a DTCP-over-SSH Connection to the Mediation Device	596
Configuring the Mediation Device to Provision Traffic Mirroring	597
Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring	597
Example: Configuring Traffic That Is Mirrored Using DTCP-Initiated Subscriber Secure Policy	598
Terminating DTCP-Initiated Subscriber Traffic Mirroring Sessions	601
Configuring DTCP Messages Used for DTCP-Initiated Subscriber Secure Policy Mirroring 	603
ADD (DTCP)	604
DELETE (DTCP)	609
ENABLE (DTCP)	611
LIST (DTCP)	613
Example: Using DTCP Messages to Trigger, Verify, and Remove Traffic Mirroring for Subscribers	615
Creating DTCP ADD Messages to Trigger Traffic Mirroring	616
Creating DTCP ENABLE Messages to Trigger Traffic Mirroring	617
Using LIST Messages to Verify That Subscriber Traffic Is Being Mirrored	618
Using DELETE Messages to Remove Traffic Mirroring Triggers	619
Verifying That Traffic Mirroring Was Stopped on the Subscriber Interfaces	620
Configuring Subscriber Secure Policy Support for IPv4 Multicast Traffic 	621
Subscriber Secure Policy Support for IPv4 Multicast Traffic	621
Triggering the Mirroring of IPv4 Multicast Traffic	621
Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic	622

Configuring Intercept-Related Information for Subscriber Secure Policy | 623

Intercept-Related Events Transmitted to the Mediation Device | 623

SNMP Traps for Subscriber Secure Policy LAES Compliance | 624

Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring | 625

Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring | 626

Remote Device and Service Management

Configuring Remote Device Services Management | 631

Remote Device Services Manager (RDSM) Overview | 631

Remote Services | 633

Process Flows for RDSM Provisioning and Deprovisioning | 633

RDSM Dictionary for Implementing Service Actions | 638

Additional Features for Use with an RDSM Access Model | 642

Response to the External Authority by authd on Success or Failure | 643

Operator Reconfiguration of Remote Devices | 645

External Notification for Service Processing ERRMSG Events | 647

Benefits of Remote Device Service Management | 647

Configuring Remote Device Management for Service Provisioning | 648

Reconfiguring a Remote Device for RDSM | 652

Reloading a Dictionary File for RDSM | 653

Configuring TCP Port Forwarding for Remote Subscriber Services | 655

TCP Port Forwarding for Remote Device Management | 655

Benefits of TCP Port Forwarding | 657

Configure TCP Port Forwarding for Remote Device Management | 659

Tracing TCP Port Forwarding Events for Troubleshooting | 662

Configuring the TCP Port Forwarding Trace Log Filename | 663

Configuring the Number and Size of TCP Port Forwarding Log Files | 663

Configuring Access to the TCP Port Forwarding Log File | 663

Configuring a Regular Expression for TCP Port Forwarding Messages to Be Logged | 664

Configuring the TCP Port Forwarding Tracing Flags | 664

Configuring the Severity Level to Filter Which TCP Port Forwarding Messages Are Logged | 665

Configuring IPFIX Mediation for Remote Device Monitoring | 667

IPFIX Mediation on the BNG | 667

Template ID Reconciliation | 669

IPFIX Mediation and Network Analytics | 670

Benefits of IPFIX Mediation | 671

Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data | 672

Collection and Export of Local Telemetry Data on the IPFIX Mediator | 677

Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector | 677

Benefits of Telemetry Data Collection | 679

Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator | 680

9

Troubleshooting

Contacting Juniper Networks Technical Support | 687

Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support | 687

10

Configuration Statements and Operational Commands

Configuration Statements | 693

access-domain (Remote Device Management) | 703

accounting (Dynamic IGMP Interface) | 705

accounting (Dynamic MLD Interface) | 706

action | 707

address (LRF Profile) | 708

address-mapping (Application Identification) | 709

adf (Dynamic Firewalls) | 710

adjustment-control-profiles | 712

adjust-minimum (Dynamic Shaping and Scheduling) | 713

adjust-percent (Dynamic Schedulers) | 714

agent (Analytics) | 715

aggregate (Hierarchical Policier) | 718

alt-name (Application Identification) | 719

analytics | 720

ancp (Adjustment Control Profiles) | 727

application (Adjustment Control Profiles) | 729

application (Application Identification) | 730

application-groups (PCC Rules) | 732

application-identification (Application Identification) | 734

application-identification-profile (Service Set) | 736

applications (PCC Rules) | 737

apply-groups (Subscriber Secure Policy) | 739

apply-groups-except (Subscriber Secure Policy) | 740

authentication-order | 741

bandwidth (Tunnel Services) | 743

bandwidth-limit (Policer) | 745

bandwidth-percent | 747

buffer-size (Dynamic Scheduling) | 750

burst-size-limit (Hierarchical Policer) | 752

burst-size-limit (Policer) | 754

bytes (Dynamic Traffic Shaping) | 756

cacheable (Application Identification) | 757

captive-portal-content-delivery (Captive Portal Content Delivery) | 758

captive-portal-content-delivery-profile (Services) | 761

cell-mode (Dynamic Traffic Shaping) | 763

chain-order (Application Identification) | 765

check-bytes (Application Identification) | 766

class (Defining Login Classes) | 767

class-of-service (Dynamic Profiles) | 778

classifiers (Dynamic CoS Application) | 780

code (Application Identification) | 781

collector (LRF Profile) | 782

collector (LRF Rule) | 783

color-aware | 784

color-blind | 786

committed-burst-size | 788

committed-information-rate | 790

compatibility (Application Identification) | 792

connection-limit | 793

context (Application Identification) | 795

delay-buffer-rate (Dynamic Traffic Shaping) | 797

description (Application Identification) | 798

destination (Application Identification) | 799

destination (LRF Profile) | 800

destination-address (Subscriber Secure Policy) | 801

destination-port (Subscriber Secure Policy) | 802

ddos-protection (DDoS) | 803

dhcp-tags (Adjustment Control Profiles) | 807

direction (Application Identification) | 809

direction (Service Data Flow Filters) | 810

disable (Dynamic IGMP) | 811

disable (Dynamic MLD) | 812

download (Application Identification) | 813

drop-policy (Subscriber Secure Policy) | 814

drop-profile (Dynamic Schedulers) | 815

drop-profile-map (Dynamic Schedulers) | 817

dscp (Dynamic Classifiers) | 818

dscp (Dynamic Rewrite Rules) | 819

dscp (Subscriber Secure Policy) | 820

dscp-ipv6 (Dynamic Classifiers) | 821

dscp-ipv6 (Dynamic Rewrite Rules) | 822

dtcp-only (System Services) | 823

dynamic-class-of-service-options (Dynamic Traffic Shaping) | 824

dynamic-profiles | 826

effective-shaping-rate | 840

enable-performance-mode (Application Identification) | 841

enhanced-mode | 842

enhanced-mode-override | 845

enhanced-policer | 846

excess-burst-size | 847

excess-priority (Dynamic Schedulers) | 849

excess-rate (Dynamic Schedulers) | 850

excess-rate (Dynamic Traffic Shaping) | 852

excess-rate-high (Dynamic Traffic Shaping) | 854

excess-rate-low (Dynamic Traffic Shaping) | 856

exclude (Dynamic MLD Interface) | 857

fail-filter (Dynamic Profiles) | 858

family (Dynamic Firewalls) | 859

family (Dynamic Standard Interface) | 861

fast-update-filter (Dynamic Firewalls) | 864

filter (Configuring) | 866

filter (Dynamic Profiles Filter Attachment) | 868

filter (Dynamic Profiles Filter Creation) | 870

filter (Dynamic Interface Unit) | 872

filter-specific | 874

firewall (Dynamic Firewalls) | 876

flow-descriptions | 879

flow-tap | 881

flow-tap-dtcp | 883

flows (PCC Rules) | 885

format (LRF Profile) | 886

forwarding-class (Dynamic Scheduler Maps) | 887

forwarding-class (PCC Action Profiles) | 888

forwarding-class (Subscriber Secure Policy) | 889

fpc (MX Series 5G Universal Routing Platforms) | 890

frame-mode (Dynamic Traffic Shaping) | 892

from (Captive Portal Content Delivery Tags) | 894

from (PCC Rules) | 895

from (Subscriber Secure Policy) | 897

gate-status | 898

group (Dynamic IGMP Interface) | 900

group (Dynamic MLD Interface) | 902

group-count (Dynamic MLD Interface) | 903

group-increment (Dynamic MLD Interface) | 904

group-limit (Dynamic IGMP Interface) | 905

group-limit (Dynamic MLD Interface) | 906

group-policy (Dynamic IGMP Interface) | 907

group-policy (Dynamic MLD Interface) | 908

guaranteed-rate (Dynamic Traffic Shaping) | 909

hierarchical-policer | 911

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) | 914

http-log-multiple-transactions (LRF Profile) | 916

icmp-mapping (Application Identification) | 917

ieee-802.1 (Dynamic Classifiers) | 918

ieee-802.1 (Dynamic Rewrite Rules) | 919

if-exceeding (Hierarchical Policer) | 920

if-exceeding (Policer) | 922

igmp (Dynamic Profiles) | 924

immediate-leave (Dynamic IGMP Interface) | 926

immediate-leave (Dynamic MLD Interface) | 927

inet (Subscriber Secure Policy) | 929

inet-precedence (Dynamic Classifiers) | 930

inet-precedence (Dynamic Rewrite Rules) | 931

inet6 (Subscriber Secure Policy) | 932

input (Dynamic Service Sets) | 933

inputs (Analytics) | 934

interface (Dynamic IGMP) | 938

interface (Dynamic Interface Sets) | 940

interface (Dynamic MLD) | 942

interface (Dynamic Routing Options) | 944

interface-service (Services Interfaces) | 945

interface-set (Dynamic Profiles) | 946

interface-shared | 948

interface-specific (Dynamic Firewalls) | 949

interfaces (Dynamic CoS Definition) | 950

interfaces (Static and Dynamic Subscribers) | 952

ip-protocol-mapping (Application Identification) | 959

ipv4-address (Steering Path) | 960

ipv6-address (Steering Path) | 961

keep-existing-steering | 962

local-port-range | 963

local-ports | 965

logging-rule (PCC Action Profile) | 967

logical-bandwidth-policer | 968

logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces) | 969

logical-interface-policer | 970

logical-system (Subscriber Secure Policy) | 973

login | 974

loss-priority (Dynamic Schedulers) | 979

loss-priority high then discard (Three-Color Policer) | 980

max-queues-per-interface | 982

match-order (Dynamic Firewalls) | 984

maximum-bit-rate (PCC Action Profiles) | 985

member (Application Identification) | 987

mld (Dynamic Profiles) | 988

multicast (Dynamic Routing Options) | 990

multicast-interception (Subscriber Secure Policy) | 991

netconf (Remote Device Management) | 992

no-accounting | 995

no-qos-adjust (Dynamic Routing Options) | 996

oif-map (Dynamic IGMP Interface) | 997

oif-map (Dynamic MLD Interface) | 998

order (Application Identification) | 999

order-priority (Application Identification) | 1000

output (Dynamic Service Sets) | 1001

outputs (Analytics) | 1002

output-traffic-control-profile (Dynamic CoS Definition) | 1006

overhead-accounting (Dynamic Traffic Shaping) | 1007

passive (Dynamic IGMP Interface) | 1008

passive (Dynamic MLD Interface) | 1009

path (Steering) | 1010

pattern (Application Identification) | 1011

pcc-action-profile (PCC Rules) | 1012

pcc-action-profiles | 1014

pcc-context | 1016

pcc-rule | 1018

pcc-rulebases (PCEF) | **1020**
pcc-rulebases (PCEF Profile) | **1022**
pcc-rules (PCEF) | **1024**
pcc-rules (PCEF Profile) | **1026**
pcef (Dynamic Profiles) | **1028**
pcef-profile (Service Set) | **1029**
peak-burst-size | **1030**
peak-information-rate | **1032**
physical-interface-policer | **1034**
policer (Configuring) | **1036**
policy (Subscriber Secure Policy) | **1038**
policy-based-logging (LRF Profile) | **1040**
policy-options (Dynamic Profiles) | **1041**
policy-statement | **1042**
port (LRF Profile) | **1048**
port-range (Application Identification) | **1049**
post-service-filter (Dynamic Service Sets) | **1050**
pppoe-tags (Adjustment Control Profiles) | **1051**
precedence | **1053**
premium (Hierarchical Policer) | **1055**
priority (Dynamic Schedulers) | **1057**
priority (Application Identification With Next Gen Services) | **1058**
profile (Access) | **1059**
profile (Captive Portal Content Delivery) | **1066**
profile (LRF) | **1068**
profile (Services PCEF) | **1070**
profiles (PCEF) | **1071**
profile-type (Dynamic Service Profiles) | **1073**
promiscuous-mode (Dynamic IGMP Interface) | **1074**
protocol (Application Identification) | **1075**
protocol (Dynamic Schedulers) | **1076**
protocol (Flow Descriptions) | **1077**
protocol (Subscriber Secure Policy) | **1078**
protocols (DDoS) | **1079**

protocols (Dynamic Profiles) | **1091**

provisioning-method (Remote Device Management) | **1094**

radius (Access Profile) | **1095**

radius-coa (Adjustment Control Profiles) | **1099**

radius-flow-tap | **1101**

radius-server | **1103**

rate-limit | **1109**

rebalance-periodic (Aggregated Ethernet Subscriber Interfaces) | **1110**

redirect (PCC Action Profiles) | **1111**

remote-address | **1113**

remote-device-management | **1115**

remote-port-range | **1117**

remote-ports | **1119**

report (LRF Rule) | **1121**

rewrite-rules (Dynamic CoS Interfaces) | **1122**

routing-engine-services | **1123**

routing-options (Dynamic Profiles) | **1124**

routing-instance (Subscriber Secure Policy) | **1126**

routing-instance (PCC Action Profiles) | **1127**

rpf-check (Dynamic Profiles) | **1129**

rule (Captive Portal Content Delivery) | **1130**

rule (LRF) | **1132**

rule-set (Captive Portal Content Delivery) | **1133**

scheduler (Dynamic Scheduler Maps) | **1134**

scheduler-map (Dynamic Traffic Shaping) | **1135**

scheduler-maps (Dynamic CoS Definition) | **1136**

schedulers (Dynamic CoS Definition) | **1137**

service (Dynamic Profiles) | **1138**

service (Dynamic Service Sets) | **1139**

service-agents (Analytics) | **1141**

service-device (Remote Device Management) | **1143**

service-filter (Dynamic Service Sets) | **1146**

service-interface (Services Interfaces) | **1147**

service-set (Application-Aware Control Policy) | **1148**

service-set (Dynamic Service Sets) | **1150**

services (Captive Portal Content Delivery) | **1152**

session-options | **1154**

shaping-rate (Dynamic Traffic Shaping and Scheduling) | **1158**

shared-name | **1160**

signature (Application Identification) | **1161**

single-rate | **1162**

snmp (Subscriber Secure Policy) | **1163**

source (Application Identification) | **1164**

source (Dynamic IGMP Interface) | **1165**

source (Dynamic MLD Interface) | **1166**

source-address (Subscriber Secure Policy) | **1167**

source-address (LRF Profile) | **1168**

source-count (Dynamic MLD Interface) | **1169**

source-increment (Dynamic MLD Interface) | **1170**

source-ipv4-address | **1171**

source-port (Subscriber Secure Policy) | **1172**

ssh (System Services) | **1173**

ssm-map (Dynamic IGMP Interface) | **1175**

ssm-map (Dynamic MLD Interface) | **1177**

ssm-map-policy (Dynamic IGMP Interface) | **1179**

ssm-map-policy (Dynamic MLD Interface) | **1181**

stacked-interface-set (Dynamic Profiles) | **1183**

static (Dynamic IGMP Interface) | **1185**

static (Dynamic MLD Interface) | **1186**

static-policy-control | **1187**

steering | **1189**

subscriber-leave-timer | **1191**

tags (Application Identification) | **1192**

targeted-distribution (Dynamic Demux Interfaces over Aggregated Ethernet) | **1193**

targeted-distribution (Static Interfaces over Aggregated Ethernet) | **1194**

tcp-forwarding (Processes) | **1195**

tcp-forwarding (Remote Device Management) | **1196**

template (LRF Profile) | **1199**

template (LRF Rule) | 1200

template-tx-interval (LRF Profile) | 1201

template-type (LRF Profile) | 1202

term (Captive Portal Content Delivery) | 1204

term (Dynamic Profiles) | 1206

then (Captive Portal Content Delivery) | 1208

then (LRF rule) | 1210

then (PCC Rules) | 1211

three-color-policer (Configuring) | 1213

time-limit (LRF Rule) | 1215

traceoptions (Analytics Agent) | 1216

traceoptions (Captive Portal Content Delivery) | 1218

traceoptions (TCP Port Forwarding) | 1220

traffic-control-profiles (Dynamic CoS Definition) | 1222

transmit-rate (Dynamic Schedulers) | 1224

trigger-type (LRF Profile) | 1226

tunnel-services (Chassis) | 1227

two-rate | 1229

type (Application Identification) | 1230

type (ICMP Mapping for Application Identification) | 1231

uid (Dynamic Profiles) | 1232

uid-reference | 1233

unit (Dynamic Profiles Standard Interface) | 1234

unit (Dynamic Traffic Shaping) | 1239

url | 1241

user (Access) | 1243

vendor-support | 1246

version (Dynamic IGMP Interface) | 1247

version (Dynamic MLD Interface) | 1248

vlan-tag (Dynamic Classifiers) | 1249

vlan-tag (Dynamic Rewrite Rules) | 1250

volume-limit (LRF Rule) | 1251

Operational Commands | 1253

clear firewall | 1256

clear igmp membership | 1259

clear interfaces statistics | 1263

clear mld membership | 1265

clear remote-device-management statistics | 1267

clear services application-identification application-system-cache | 1269

clear services application-identification statistics | 1270

clear services captive-portal-content-delivery statistics | 1273

clear services lrf collector statistics | 1275

clear services lrf statistics | 1276

clear tcp-forwarding connections | 1277

clear tcp-forwarding statistics | 1280

request interface rebalance (Aggregated Ethernet for Subscriber Management) | 1284

request network-access aaa subscriber add session-id | 1285

request network-access aaa subscriber delete session-id | 1287

request network-access aaa subscriber modify session-id | 1290

request network-access aaa subscriber set session-id | 1292

request services application-identification application | 1294

request services application-identification download | 1296

request services application-identification download status | 1297

request services application-identification group | 1298

request services application-identification install | 1300

request services application-identification install status | 1302

request services application-identification proto-bundle-status | 1303

request services application-identification uninstall | 1304

request services application-identification uninstall status | 1305

request services remote-device-management reconfigure service-device | 1306

request services remote-device-management reload-dictionary | 1308

show class-of-service | 1310

show class-of-service adjustment-control-profile | 1313

show class-of-service interface | 1315

show class-of-service interface-set | **1355**

show class-of-service scheduler-hierarchy interface | **1358**

show class-of-service scheduler-hierarchy interface-set | **1361**

show class-of-service scheduler-map | **1363**

show class-of-service traffic-control-profile | **1367**

show dynamic-profile session | **1372**

show firewall | **1378**

show firewall log | **1389**

show firewall templates-in-use | **1393**

show igmp group | **1395**

show igmp interface | **1400**

show interfaces statistics | **1405**

show interfaces targeting (Aggregated Ethernet for Subscriber Management) | **1422**

show mld group | **1424**

show mld interface | **1429**

show network-access aaa subscribers session-id | **1434**

show services analytics agent | **1442**

show remote-device-management service-devices | **1445**

show remote-device-management services | **1453**

show remote-device-management statistics | **1456**

show remote-device-management subscribers | **1461**

show remote-device-management summary | **1465**

show services application-identification application | **1469**

show services application-identification application-system-cache | **1477**

show services application-identification commit-status (Next Gen Services) | **1482**

show services application-identification counter | **1484**

show services application-identification group | **1487**

show services application-identification statistics application-groups | **1492**

show services application-identification statistics applications | **1494**

show services application-identification status | **1496**

show services application-identification version | **1499**

show services captive-portal-content-delivery | **1500**

show services lrf collector statistics | **1506**

show services lrf rule statistics | **1508**

show services lrf statistics | **1510**
show services lrf template | **1512**
show services pcef pic | **1515**
show services pcef subscribers | **1517**
show services service-sets summary | **1525**
show subscribers | **1527**
show subscribers summary | **1576**
show tcp-forwarding status | **1585**

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | xxxiii
- Using the Examples in This Manual | xxxiii
- Documentation Conventions | xxxv
- Documentation Feedback | xxxviii
- Requesting Technical Support | xxxviii

Use this guide to understand conceptual and configuration information about dynamic class of service, policy filters, and traffic policing; dynamic IGMP and MLD for access to multicast networks; application-aware policy control; HTTP redirect services to capture subscriber network requests and send them to a captive portal for authentication and access to authorized Web resources; and subscriber secure policy traffic mirroring to mirror subscriber traffic and monitor events related to the mirrored session.

Documentation and Release Notes

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

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

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

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

For more information about the **load** command, see [CLI Explorer](#).

Documentation Conventions

[Table 1 on page xxxvi](#) defines notice icons used in this guide.

Table 1: Notice Icons







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

Table 2 on page xxxvi 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
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>

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

Convention	Description	Examples
<i>Italic text like this</i>	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 <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i> >;
(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 (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions


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

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

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



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

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

Creating a Service Request with JTAC

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

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

1

PART

Subscriber Service Activation and Management

Subscriber Service Activation and Management | 3

Subscriber Service Activation and Management

IN THIS CHAPTER

- [Dynamic Service Management with RADIUS | 3](#)
- [Service Activation and Deactivation Using the CLI Instead of RADIUS | 20](#)
- [Management of Subscriber Services with Multiple Instances | 25](#)

Dynamic Service Management with RADIUS

IN THIS SECTION

- [Using RADIUS Dynamic Requests for Subscriber Access Management | 3](#)
- [Configuring RADIUS-Initiated Dynamic Request Support | 5](#)
- [RADIUS-Initiated Change of Authorization \(CoA\) Overview | 6](#)
- [RADIUS-Initiated Disconnect Overview | 10](#)
- [Usage Thresholds for Subscriber Services | 11](#)
- [Subscriber Session Logins and Service Activation Failures Overview | 12](#)
- [Configuring How Service Activation Failures Affect Subscriber Login | 18](#)
- [Error-Cause Codes \(RADIUS Attribute 101\) for Dynamic Requests | 19](#)
- [Verifying RADIUS Dynamic-Request Statistics | 20](#)

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.

By default, the router monitors UDP port 3799 for CoA requests from RADIUS servers. You can also configure a nondefault port for RADIUS servers. You must either use the default port for all RADIUS servers or configure the same nondefault port for all RADIUS servers. This rule applies at both the global access and access profile levels.

NOTE: Any other configuration results in a commit check failure. Multiple port numbers—that is, different port numbers for different servers—are not supported.

Benefits of Radius Dynamic Requests

Enables simplified central management of subscriber sessions by sending unsolicited changes to subscriber sessions, including changes in attributes, service activation, service deactivation, and session termination.

SEE ALSO

[RADIUS-Initiated Change of Authorization \(CoA\) Overview | 6](#)

[RADIUS-Initiated Disconnect Overview | 10](#)

[Configuring RADIUS-Initiated Dynamic Request Support | 5](#)

[RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework](#)

[Error-Cause Codes \(RADIUS Attribute 101\) for Dynamic Requests | 19](#)

Configuring RADIUS-Initiated Dynamic Request Support

The router uses the list of specified RADIUS authentication servers for both authentication and dynamic request operations. By default, the router monitors UDP port 3799 for dynamic requests, also known as Change of Authorization (CoA) 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
```

When you configure dynamic request ports, you must do one of the following:

- Use the default port for all RADIUS servers at both the global access level and in all access profiles.
- Configure the same nondefault port for all servers at both the global access level and in all access profiles.

NOTE: Any other configuration results in a commit check failure. Multiple port numbers—that is, different port numbers for different servers—are not supported.

To specify a global dynamic request port:

```
[edit access]
user@host# set radius-server server-address dynamic-request-port port-number
```

To specify the dynamic request port for a specific access profile:

```
[edit access]
user@host# set profile profile-name radius-server server-address dynamic-request-port port-number
```

Consider the following scenarios:

- The following configuration uses the default port for both a server globally and a different server in the access profile. This is a valid configuration.

```
[edit access]
user@host# set radius-server 192.0.2.1
user@host# set profile ap1 radius-server 192.0.2.3
```

- The following configuration specifies nondefault port 50201 for both a server globally and a different server in the access profile. This is a valid configuration.

```
[edit access]
user@host# set radius-server 192.0.2.1 dynamic-request-port 50201
user@host# set profile ap1 radius-server 192.0.2.3 dynamic-request-port 50201
```

- The following configuration specifies port 50201 globally for a server and port 51133 for the same server in the ap1 access profile. This is an invalid configuration and commit check fails, because multiple nondefault ports are not supported.

```
[edit access]
user@host# set radius-server 192.0.2.1 dynamic-request-port 50201
user@host# set profile ap1 radius-server 192.0.2.1 dynamic-request-port 51133
```

- The following configuration uses the default port 3799 for one server globally and port 51133 for another server globally. This is an invalid configuration and the commit check fails, because for all servers you must configure either the default port or the same nondefault port.

```
[edit access]
user@host# set radius-server 192.0.2.1
user@host# set radius-server 192.0.2.3 dynamic-request-port 51133
```

SEE ALSO

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

RADIUS-Initiated Change of Authorization (CoA) Overview

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

CoA Messages

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

- CoA-Request (43)
- 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 3 on page 7 shows the identification attributes for CoA operations.

Table 3: Identification Attributes

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

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

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

Table 4 on page 8 shows the session attributes for CoA operations. Any additional client attributes that you include depend on your particular session requirements.

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

Message Exchange

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

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

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

NOTE: The AAA Service Framework processes one dynamic request at a time per subscriber. If the framework receives a second dynamic request (either another CoA or a Disconnect-Request) while processing a previous request for the same subscriber, the framework responds with a CoA-NAK message. Starting in Junos OS Release 15.1R5, CoA-Request retry messages are ignored and no CoA-NAK is sent in response to them.

Bulk CoA Transactions

Starting in Junos OS Release 17.2R1, bulk CoA requests are supported to improve the processing efficiency of RADIUS-based subscriber services on the BNG. The bulk CoA functionality enables the accumulation of a series of CoA requests and commits all of them together, in bulk, automatically.

Bulk CoA transactions are particularly valuable for business services. RADIUS-based subscriber services use the Juniper Networks VSAs, Activate-Service (26-65) and Deactivate-Service (26-66). The VSAs are provided in RADIUS-Accept messages during login or in CoA requests after login.

For conventional, dynamic service profile-based services, up to 12 service activations can easily fit within either RADIUS message. However, the op-script based services used by businesses typically have scaling requirements that exceed the capacity of either message. This means that specifying and activating all the services needed in a given subscriber session may require using an Accept-Access message and multiple CoA requests.

Each CoA request message is independent of previous and future CoA requests in the same subscriber session. All service-activations and deactivations in a message are processed before a CoA response is offered. The CoA request provides a way to incrementally modify a subscriber session without affecting existing services that are already present.

For op-script based services, the service sessions are collaboratively created by the authd and essmd processes such that the last operation initiates a commit to apply all resultant static business logical interfaces from the CoA request. Because the commit time is generally the largest part of applying a static business service, there is an advantage to packing as many service-activations or deactivations as will fit within a RADIUS message to efficiently use the commit window. Until the commit operation completes, the BNG cannot accept a subsequent CoA request to apply additional business services for the same subscriber session.

Bulk CoA improves the efficiency of commit processing by using a single commit action for all services in the bulk transaction. The bulk transaction has the effect of managing a series of requests as a single meta-request. It defers the commit processing until the final CoA request in the bulk transaction is received.

Bulk CoA requires each individual request to contain a single instance of the Juniper Networks Bulk-CoA-Transaction-Id VSA (26-194). This VSA identifies requests as belonging to the same bulk transaction; 26-194 must have the same value in all CoA requests in the bulk series. Each successive bulk transaction in the session must have a different identifier; for example, three successive bulk transactions can have IDs of 1, 2, and 1, but cannot have successive IDs of 1, 1, and 2. In practice, the Bulk-CoA-Transaction-Id value typically is incremented for multiple bulk transactions, but this is not required. An ID value used in a given subscriber session can also be used in different subscriber sessions.

Each CoA request within a bulk transaction has its own unique identifier, provided by a single instance of the Bulk-CoA-Identifier VSA (26-195) in each CoA. An increasing series of values for the ID is typical but not enforced. Values can be reused within a given session and between sessions. The final CoA request in the series is identified by having a value of 0xFFFFFFFF for the Bulk-CoA-Identifier.

Benefits of Radius-Initiated Change of Authorization

Enables changes in attribute values to be dynamically pushed to subscriber sessions, as well as dynamic activation and deactivation of subscriber services.

SEE ALSO

[Using RADIUS Dynamic Requests for Subscriber Access Management | 3](#)

[RADIUS-Initiated Disconnect Overview | 10](#)

[Configuring RADIUS-Initiated Dynamic Request Support | 5](#)

[Usage Thresholds for Subscriber Services | 11](#)

[RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework](#)

RADIUS-Initiated Disconnect Overview

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

Disconnect Messages

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

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

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

Qualifications for Disconnect

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

Message Exchange

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

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.

Benefits of Radius-Initiated Disconnects

Enables a RADIUS server to dynamically terminate subscriber sessions. This centralized subscriber management feature simplifies handling large numbers of subscribers because operator termination would otherwise require action on the router.

SEE ALSO

[Using RADIUS Dynamic Requests for Subscriber Access Management | 3](#)

[Configuring RADIUS-Initiated Dynamic Request Support | 5](#)

Usage Thresholds for Subscriber Services

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

Subscriber management supports two types of usage thresholds—traffic volume and time. You use Juniper Networks VSAs to set the usage thresholds. The VSAs are transmitted in RADIUS Access-Accept messages for dynamically activated services, or in RADIUS-initiated CoA-Request messages for existing services. The volume threshold sets the maximum amount of the total input and output traffic that can use the service before the service is deactivated. A time threshold sets the maximum length of time that the service can be active. [Table 5 on page 12](#) shows the VSAs used for volume and time thresholds.

Table 5: Juniper Network VSAs Used for Service Thresholds

Attribute Number	Attribute Name	Description	Value
26-67	Service-Volume	Amount of input and output traffic, in MB, that can use the service; service is deactivated when the volume is exceeded. Tagged VSA, which supports 8 tags (1-8). The router polls the traffic in 10 minute intervals.	<ul style="list-style-type: none"> • range = 0 through 16777215 MB • 0 = no limit
26-68	Service-Timeout	Number of seconds that the service can be active; service is deactivated when the timeout expires. Tagged VSA, which supports 8 tags (1-8).	<ul style="list-style-type: none"> • range = 0 through 16777215 seconds • 0 = no timeout
26-179	Service-Volume-4GB	Amount of input and output traffic, in 4GB units, that can use the service; service is deactivated when the volume is exceeded. Tagged VSA, which supports 8 tags (1-8). The router polls the traffic in 10 minute intervals.	<ul style="list-style-type: none"> • range = 0 through 16777215 4GB units • 0 = no limit
26-180	Update-Service	New values of service and time quotas for an existing service. Tagged VSA, which supports 8 tags (1-8).	string: <i>service-name</i>

SEE ALSO

[RADIUS-Initiated Change of Authorization \(CoA\) Overview](#) | 6

Subscriber Session Logins and Service Activation Failures Overview

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

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

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

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

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

- **required-at-login**—Activation is required. Failure for any reason causes the Network-Family-Activate-Request for that network family to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber. This is the default behavior for the **dynamic-profile** service type.
- **optional-at-login**—Activation is optional. Failure due to configuration errors does not prevent activation of the address family; it allows subscriber access. Failure for any other reason causes network family activation to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber. This is the default behavior for the **extensible-service** service type.

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

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

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

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

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

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

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

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

Service and Network Family Activation Process

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

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

- When the service activation fails due to a configuration error and the service is optional:
 - a. authd deletes the cached service activation attributes for the service.

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

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

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

- b. authd sends a NACK in response to the family activation, which causes the client application to terminate the subscriber's login.
- When the service activation fails due to a nonconfiguration error and the service is either optional or required:
 - a. authd deletes the cached service activation attributes for the service.

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

- b. authd sends a NACK in response to the family activation, which causes the client application to terminate the subscriber's login.
- When the service activation succeeds:
 - a. authd activates the service.

- b. authd sends an ACK in response to the family activation request and the family is activated.
 - c. The subscriber login proceeds.
5. Unless service activation was required and failed, causing the family activation to fail in the first request, the client application may send a second request, but only for the family not requested the first time. If the first request was for IPv4, then the second request can only be for IPv6. If the first request was for IPv6, then the second request can only be for IPv4.
 6. authd reviews the cached service activation attributes and requests activation for the services associated with the requested address family.
 7. What authd does next depends on whether the service activation request fails and whether the service is optional or required.
 - When the service activation fails due to a configuration error and the service is optional:
 - a. authd deletes the cached service activation attributes for the service.

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

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

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

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

- If the first family activation was successful and that subscriber logged in, failure of the second request does not halt the current subscriber login. This event also does not cause authd to log out the previous (first request) subscriber.
- If the first family activation was unsuccessful, failure of the second request causes the client application to terminate the current subscriber login.
- When the service activation fails due to a nonconfiguration error and the service is either optional or required:
 - a. authd deletes the cached service activation attributes for the service.

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

- b. authd sends a NACK in response to the family activation, which causes the client application to terminate the subscriber's login.
- When the service activation succeeds:
 - a. authd activates the service.
 - b. authd sends an ACK in response to the family activation request and the family is activated.
 - c. The subscriber login proceeds.

Configuring How Service Activation Failures Affect Subscriber Login

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

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

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

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

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

- Specify that service activation is optional.

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

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

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

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

- Specify that service activation is required.

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

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

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

SEE ALSO

[Subscriber Session Logins and Service Activation Failures Overview | 12](#)

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 6 on page 19](#) describes the error-cause codes.

Table 6: Error-Cause Codes (RADIUS Attribute 101)

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

SEE ALSO

[RADIUS-Initiated Change of Authorization \(CoA\) Overview | 6](#)
[RADIUS-Initiated Disconnect Overview | 10](#)

Verifying RADIUS Dynamic-Request Statistics

Purpose

Display RADIUS dynamic request statistics and information.

Action

- To display RADIUS dynamic request statistics:

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

SEE ALSO

| [CLI Explorer](#)

Service Activation and Deactivation Using the CLI Instead of RADIUS

IN THIS SECTION

- [CLI-Activated Subscriber Services | 20](#)
- [Local and Remote Service Activation and Deactivation Using the CLI | 21](#)

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 then log in again to complete the change. For example, a service provider might allow subscribers to log in and initially use the default service, which provides basic features. After the default service is established, the provider might then use CLI-activation to upgrade qualified subscribers to an advanced service, in addition to retaining the initial 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*.

SEE ALSO

[Local and Remote Service Activation and Deactivation Using the CLI | 21](#)

[Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers | 50](#)

[Default Subscriber Service Overview](#)

Local and Remote Service Activation and Deactivation Using 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. This method is an alternative to using external actions by your RADIUS server.

Starting in Junos OS Release 18.3R1, when the dynamic service profile is configured with the **profile-type remote-device-service** statement, the CLI statements trigger the remote device services manager (RDSM) to provision or deprovision the service on a remote device.

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*.

However, this caveat does not apply if the service was provisioned on a remote device by the RDSM at the request of PCRF as the external authority supplying the service information. When you issue the command to activate or deactivate such a service, RDSM handles the service action.

To use the CLI to activate a subscriber service:

1. (Optional) Verify the subscriber's ID, and ensure that provisioning is not enabled. To display the session IDs of all current subscribers, use the **show subscribers detail** or **show network-access aaa subscribers** command.

```
user@host> show network-access aaa subscribers session-id 55 detail
```

```
Type: dhcp
Username: user23@example.net
Stripped username: user23
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:retail-onlinecompany-ca
Access-profile:retailer-onlinecompany-sjc
Session ID: 55
Accounting Session ID: 55
Multi Accounting Session ID: 0
IP Address: 192.168.44.104
Authentication State: AuthStateActive
Accounting State: Acc-Start-Send
Provisioning-type: none
Service name: basic-service
Service State: SvcActive
```



```
Session ID: 56
Session uptime: 00:01:45
```

2. Activate the service for the subscriber.

```
user@host> request network-access aaa subscriber add session-id 55 service-profile gold-service
```

3. (Optional) Verify that the new service is activated for the subscriber. (The initial **basic-service** is also listed because it has not been deactivated.)

```
user@host> show network-access aaa subscribers session-id 55 detail
```

```
Type: dhcp
Username: user23@example.net
Stripped username: user23
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:retail-onlinecompany-ca
Access-profile:retailer-onlinecompany-sjc
Session ID: 55
Accounting Session ID: 55
Multi Accounting Session ID: 0
IP Address: 192.168.44.104
Authentication State: AuthStateActive
Accounting State: Acc-Start-Send
Provisioning-type: none
Service name: basic-service
  Service State: SvcActive
  Session ID: 56
  Session uptime: 00:02:15
Service name: gold-service
  Service State: SvcActive
  Session ID: 57
  Session uptime: 00:00:30
```

To use the CLI to deactivate a subscriber service:

1. Display the active services for the specified subscriber. The following example shows that the **basic-service** and **gold-service** are active.

```
user@host> show network-access aaa subscribers session-id 55 detail
```

```

Type: dhcp
Username: user23@example.net
Stripped username: user23
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:retail-onlinecompany-ca
Access-profile:retailer-onlinecompany-sjc
Session ID: 55
Accounting Session ID: 55
Multi Accounting Session ID: 0
IP Address: 192.168.44.104
Authentication State: AuthStateActive
Accounting State: Acc-Start-Send
Provisioning-type: none
Service name: basic-service
  Service State: SvcActive
  Session ID: 56
  Session uptime: 00:02:15
Service name: gold-service
  Service State: SvcActive
  Session ID: 57
  Session uptime: 00:00:30

```

2. Deactivate the service for the subscriber. The following example deletes the subscriber's **basic-service** service.

```

user@host> request network-access aaa subscriber delete session-id 55 service-profile basic-service

```

3. (Optional) Verify that the deleted service is no longer active for the subscriber. (The **gold-service** is still listed because it has not been deactivated.)

```

user@host> show network-access aaa subscribers session-id 55 detail

```

```

Type: dhcp
Username: user23@example.net
Stripped username: user23
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:retail-onlinecompany-ca
Access-profile:retailer-onlinecompany-sjc
Session ID: 55
Accounting Session ID: 55
Multi Accounting Session ID: 0
IP Address: 192.168.44.104
Authentication State: AuthStateActive

```

```
Accounting State: Acc-Start-Send
Provisioning-type: none
Service name: gold-service
Service State: SvcActive
Session ID: 57
Session uptime: 00:00:30
```

The following sample commands illustrate CLI activation and deactivation for remote services applied by RDSM to remote devices.

- **user@host> request network-access aaa subscriber add session-id 131 service-profile "upstreamBandwidth(100,100,100)"**

```
Successful completion
```

- **user@host> request network-access aaa subscriber delete session-id 131 service-profile "upstreamBandwidth(100,100,100)"**

```
Successful completion
```

SEE ALSO

[CLI-Activated Subscriber Services | 20](#)

[Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers | 50](#)

[Default Subscriber Service Overview](#)

[Configuring Remote Device Management for Service Provisioning | 648](#)

[Remote Device Services Manager \(RDSM\) Overview | 631](#)

Management of Subscriber Services with Multiple Instances

IN THIS SECTION

- [Subscriber Services with Multiple Instances Overview | 26](#)
- [Deactivating a Single Instance of a Subscriber Service | 28](#)

- [Deactivating All Instances of a Subscriber Service | 31](#)
- [Verifying Subscriber Services with Multiple Instances | 34](#)

Subscriber Services with Multiple Instances Overview

IN THIS SECTION

- [Subscriber Service Instances and Service Parameters | 26](#)
- [CLI Deactivation of Subscriber Services with Multiple Instances | 26](#)
- [Subscriber Services with Multiple Instances in RADIUS Accounting Messages | 27](#)

Services are activated for subscribers either at login, or by using Change of Authorization (CoA) RADIUS messages or command-line interface (CLI) requests. A subscriber can have multiple instances of the same named service, provided that each instance of the subscriber service has a different set of parameters. Support for multiple instances of a subscriber service enables you to use service parameters to customize the same service to meet different needs for a particular subscriber.

Subscriber Service Instances and Service Parameters

In a subscriber access network, each subscriber has its own set of services. You can configure a specific *service instance* for a particular subscriber by specifying a *service name*, also referred to as a *service profile*, and unique service parameters for that service instance. *Service parameters* can include a combination of policy lists, filters, rate-limit profiles, class of service (CoS) profiles, and interface profiles.

For example, `filter-service(up-filter,down-filter)` and `filter-service(upstream-filter,downstream-filter)` are considered two different instances of the same service (`filter-service`) because their parameters, enclosed in parentheses after the service name, are different.

Each service instance is uniquely identified by the combination of its service name and service parameters. In CoA messages, the router identifies a subscriber service by its complete activation string, which consists of the service name and, if configured, one or more service parameters in the order specified.

CLI Deactivation of Subscriber Services with Multiple Instances

You can use the Junos OS CLI to deactivate subscriber services with multiple instances in either of the following ways:

- Deactivate a single instance of a subscriber service by specifying the name and parameters of the service to be deactivated.

With this feature, you can deactivate a particular instance of a subscriber service while other instances of that same service remain active. For example, assume that a subscriber identified by a particular session ID has two instances of filter-service activated: filter-service(up-filter,down-filter) and filter-service(upstream-filter,downstream-filter). If you specify “filter-service(up-filter,down-filter)” in the **request network-access aaa subscriber delete session-id** command, the router deactivates only filter-service(up-filter,down-filter); filter-service(upstream-filter,downstream-filter) remains active.

The ability to use both service names and service parameters to identify the particular service instance to be deactivated is analogous to the subscriber service deactivation feature in use on Juniper Networks E Series Broadband Services Routers that run JunosE Software.

- Deactivate all instances of a subscriber service by specifying only the name of the service to be deactivated, with no service parameters.

With this feature, you can deactivate all instances of the same subscriber service with a single operational command. Using the same subscriber service example, if you specify “filter-service” in the **request network-access aaa subscriber delete session-id** command, the router deactivates both filter-service(up-filter,down-filter) and filter-service(upstream-filter,downstream-filter).

Subscriber Services with Multiple Instances in RADIUS Accounting Messages

RADIUS Acct-Start, Interim-Acct, and Acct-Stop accounting messages include the subscriber service name and, if configured, service parameters. If RADIUS logging is enabled, the router logs all subscriber service attributes, including service names and parameters, in messages sent to and received from the RADIUS authentication server.

For example, assume that the router receives the following RADIUS Access-Accept message from the RADIUS server:

```
Jul 13 12:37:02 radius-access-accept: Activate-Service (Juniper-ERX-VSA) received:
Tag (1) filter-service(up-filter,down-filter)
```

[Table 7 on page 27](#) shows sample logged RADIUS Acct-Start, Interim-Acct, and Acct-Stop messages that the router sends to the RADIUS server in response to the Access-Accept message. In each of these accounting messages, the Activate-Service-Session-Name is the full activation string that includes both the service name (filter-service) and service parameters (up-filter,down-filter) to identify the service instance.

Table 7: Subscriber Services and Service Parameters in RADIUS Accounting Messages

RADIUS Accounting Message Type	RADIUS Accounting Message Text
Acct-Start	<pre>Jul 13 12:37:02 radius-acct-start: Activate-Service-Session-Name (Juniper-ERX-VSA) added: filter-service(up-filter,down-filter)</pre>

Table 7: Subscriber Services and Service Parameters in RADIUS Accounting Messages (*continued*)

RADIUS Accounting Message Type	RADIUS Accounting Message Text
Interim-Acct	Jul 13 12:47:00 radius-acct-interim: Activate-Service-Session-Name (Juniper-ERX-VSA) added: filter-service(up-filter,down-filter)
Acct-Stop	Jul 13 12:53:59 radius-acct-stop: Activate-Service-Session-Name (Juniper-ERX-VSA) added: filter-service(up-filter,down-filter)

SEE ALSO

[Deactivating a Single Instance of a Subscriber Service | 28](#)
[Deactivating All Instances of a Subscriber Service | 31](#)
[Verifying Subscriber Services with Multiple Instances | 34](#)
Deactivating a Single Instance of a Subscriber Service

For subscriber services that have multiple instances, you can use the Junos OS CLI to deactivate a service in either of the following ways:

- Deactivate a single instance of a service by specifying the name and parameters of the service to be deactivated.
- Deactivate all instances of a service by specifying only the name of the service to be deactivated.

This topic describes how to deactivate a single instance of a subscriber service.

To use the Junos OS CLI to deactivate a single instance of a subscriber service with multiple instances:

1. Display the active services for the subscriber identified by the specified session ID.

```
user@host> show network-access aaa subscribers session-id subscriber-session-id detail
```

For example, the following command displays the active services for the DHCP subscriber identified by session ID 6. In this example, two instances of economy-service are active:

economy-service(up-filter,down-filter) and **economy-service(upstrm-filter,dwnstrm-filter)**. A single instance of premium-service named **premium-service(up-filter,down-filter)** is also active.

```
user@host> show network-access aaa subscribers session-id 6 detail
```

```

Type: dhcp
Stripped username: fran2
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: attr_test_profile1
Session ID: 6
Accounting Session ID: 6
Multi Accounting Session ID: 0
IP Address: 198.51.100.10
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
Service name: economy-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 7
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:7-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: economy-service(upstrm-filter,dwnstrm-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 8
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:8-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: premium-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 9
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:9-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600

```

2. Deactivate the specified instance of a subscriber service by specifying its service name and parameters.

```
user@host> request network-access aaa subscriber delete session-id subscriber-session-id
service-profile "profile-name(parameters)"
```

For example, the following command deactivates only the instance of economy-service named economy-service(up-filter,down-filter).

```
user@host> request network-access aaa subscriber delete session-id 6 service-profile
"economy-service(up-filter,down-filter)"
```

3. (Optional) Verify that the deactivated service instance is no longer active for the subscriber.

```
user@host> show network-access aaa subscribers session-id subscriber-session-id detail
```

For example, the following command displays the services still active for the DHCP subscriber identified by session ID 6. In this example, **economy-service(up-filter,down-filter)** is no longer listed because it was deactivated, but **economy-service(upstrm-filter,dwnstrm-filter)** and **premium-service(up-filter,down-filter)** are still active.

```
user@host> show network-access aaa subscribers session-id 6 detail
```

```
Type: dhcp
Stripped username: fran2
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: attr_test_profile1
Session ID: 6
Accounting Session ID: 6
Multi Accounting Session ID: 0
IP Address: 198.51.100.13.10
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
Service name: economy-service(upstrm-filter,dwnstrm-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 8
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:8-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: premium-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
```



```

Service Activation Source: Radius
Session ID: 9
Session uptime: 00:04:36
Accounting status: on/volume+time
Service accounting session ID: 6:9-1354811427
Service accounting state: Acc-Start-Sent
Accounting interim interval: 600

```

SEE ALSO

[Deactivating All Instances of a Subscriber Service | 31](#)

[Verifying Subscriber Services with Multiple Instances | 34](#)

[Subscriber Services with Multiple Instances Overview | 26](#)

Deactivating All Instances of a Subscriber Service

For subscriber services that have multiple instances, you can use the Junos OS CLI to deactivate a service in either of the following ways:

- Deactivate a single instance of a service by specifying the name and parameters of the service to be deactivated.
- Deactivate all instances of a service by specifying only the name of the service to be deactivated.

This topic describes how to deactivate all instances of a subscriber service.

To use the Junos OS CLI to deactivate all instances of a subscriber service with multiple instances:

1. Display the active services for the subscriber identified by the specified session ID.

```
user@host> show network-access aaa subscribers session-id subscriber-session-id detail
```

For example, the following command displays the active services for the DHCP subscriber identified by session ID 6. In this example, two instances of economy-service are active:

economy-service(up-filter,down-filter) and **economy-service(upstrm-filter,dwnstrm-filter)**. A single instance of premium-service named **premium-service(up-filter,down-filter)** is also active.

```
user@host> show network-access aaa subscribers session-id 6 detail
```

```

Type: dhcp
Stripped username: fran2
AAA Logical system/Routing instance: default:default

```

```

Target Logical system/Routing instance: default:default
Access-profile: attr_test_profile1
Session ID: 6
Accounting Session ID: 6
Multi Accounting Session ID: 0
IP Address: 198.51.100.10
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
Service name: economy-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 7
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:7-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: economy-service(upstrm-filter,dwnstrm-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 8
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:8-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: premium-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 9
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:9-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600

```

2. Deactivate all instances of the specified service by specifying the service name without parameters.

```

user@host> request network-access aaa subscriber delete session-id subscriber-session-id
service-profile "profile-name"

```

For example, the following command deactivates both instances of economy-service.

```
user@host> request network-access aaa subscriber delete session-id 6 service-profile
"economy-service"
```

3. (Optional) Verify that all instances of the deactivated service are no longer active for the subscriber.

```
user@host> show network-access aaa subscribers session-id subscriber-session-id detail
```

In the following example, only **premium-service(up-filter,down-filter)** is still active. Neither **economy-service(up-filter,down-filter)** nor **economy-service(upstrm-filter,dwnstrm-filter)** is listed because all instances of economy-service were deactivated.

```
user@host> show network-access aaa subscribers session-id 6 detail
```

```
Type: dhcp
Stripped username: fran2
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: attr_test_profile1
Session ID: 6
Accounting Session ID: 6
Multi Accounting Session ID: 0
IP Address: 198.51.100.10
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
Service name: premium-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 9
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:9-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
```

SEE ALSO

[Deactivating a Single Instance of a Subscriber Service | 28](#)

[Verifying Subscriber Services with Multiple Instances | 34](#)

Verifying Subscriber Services with Multiple Instances

Purpose

Display information about the active services for a subscriber identified by the specified session ID.

Action

The following example displays information about the active services for the DHCP subscriber identified by session ID 6.

```
user@host> show network-access aaa subscribers session-id 6 detail
```

```
Type: dhcp
Stripped username: fms2
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: attr_test_profile1
Session ID: 6
Accounting Session ID: 6
Multi Accounting Session ID: 0
IP Address: 198.51.100.10
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
Service name: economy-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 7
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:7-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: economy-service(upstrm-filter,dwnstrm-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 8
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:8-1354811427
  Service accounting state: Acc-Start-Sent
```

```

Accounting interim interval: 600
Service name: premium-service
  Service State: SvcActive
Service Family: inet
Service Activation Source: Radius
Session ID: 9
Session uptime: 00:04:36
Accounting status: on/volume+time
Service accounting session ID: 6:9-1354811427
Service accounting state: Acc-Start-Sent
Accounting interim interval: 600

```

Meaning

If parameters are configured when a subscriber service with multiple instances is activated, the **Service name** field in the **show network-access aaa subscribers session-id** command displays both the service name and, in parentheses following the service name, the service parameters. If parameters are not configured for a particular service, the **show network-access aaa subscribers session-id** command displays only the service name. The value **SvcActive** in the **Service State** field indicates that the service is active.

In this example, two instances of economy-service are active: **economy-service(up-filter,down-filter)** and **economy-service(upstrm-filter,dwnstrm-filter)**. For **premium-service**, which is also active, the command output displays only the service name, indicating that no parameters were configured for this service.

SEE ALSO

[Deactivating a Single Instance of a Subscriber Service | 28](#)

[Deactivating All Instances of a Subscriber Service | 31](#)

2

PART

Configuring Dynamic Class of Service

CoS for Subscriber Access and Interfaces Overview | **39**

Configuring Scheduling and Shaping for Subscriber Access | **47**

Configuring Hierarchical CoS Scheduling on MPLS Ethernet Pseudowire Subscriber Interfaces | **67**

Allocating Dedicated Queues for Each Logical Interface Using Per-Unit Scheduling | **85**

Configuring Dedicated Queue Scaling with Hierarchical CoS or Per-Unit Scheduling | **101**

Shaping Downstream Traffic Based on Frames or Cells | **109**

Applying CoS to Households or Individual Subscribers Using ACI-Based Dynamic VLANs | **127**

Applying CoS to Households or Individual Subscribers Using Access Line Identifier Dynamic VLANs | **137**

Managing Excess Bandwidth Distribution and Traffic Bursts | **147**

Applying CoS Using Parameters Received from RADIUS | **155**

Modifying a Subscriber's Shaping Characteristics After a Subscriber is Instantiated | **183**

Applying CoS to Groups of Subscriber Interfaces | **189**

Applying CoS to Subscriber Interfaces | **217**

CoS for Subscriber Access and Interfaces Overview

IN THIS CHAPTER

- CoS for Subscriber Access Overview | 39
- Guidelines for Configuring Dynamic CoS for Subscriber Access | 40
- CoS for Aggregated Ethernet Subscriber Interfaces Overview | 44
- CoS for PPPoE Subscriber Interfaces Overview | 45

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 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 4, 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 Hierarchical CoS for Subscriber Interfaces

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

Configuring Static Hierarchical Scheduling in a Dynamic Profile

[Configuring Per-Unit Scheduling in a Dynamic Profile | 86](#)

Guidelines for Configuring Dynamic CoS for Subscriber Access

This topic describes the guidelines for configuring dynamic CoS in a subscriber access environment.

Configuration Guidelines for Hierarchical CoS and Per-Unit Scheduling

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.

Junos software supports either per-unit scheduling or hierarchical scheduling on an interface. You cannot run both types of scheduling at the same time. If CoS is active on an interface, and you change the type of scheduling configured on the interface, all traffic is dropped upon egress from the interface.

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 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 8 on page 42](#) lists the classification rule configuration for an IP demux subscriber interface with a VLAN underlying interface.

Table 8: 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 9 on page 42](#) lists the rewrite rule configuration for an IP demux subscriber interface with a VLAN underlying interface.

Table 9: 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
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 10 on page 43](#) lists the classification rule configuration for an L2TP LAC subscriber interface with a VLAN underlying interface.

Table 10: 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 11 on page 43](#) lists the rewrite rule configuration for an L2TP LAC subscriber interface with a VLAN underlying interface.

Table 11: 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

RELATED DOCUMENTATION

[CoS for Subscriber Access Overview | 39](#)

Understanding Hierarchical CoS for Subscriber Interfaces

Configuring Static Hierarchical Scheduling in a Dynamic Profile

[Configuring Per-Unit Scheduling in a Dynamic Profile | 86](#)

Configuring Static CoS for an L2TP LNS Inline Service

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.

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.
 - 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.

BEST PRACTICE: While subscribers are active on aggregated Ethernet physical interfaces with targeted distribution, we recommend that you do not change any attribute of the physical interfaces, such as MTU. Instead, perform the following steps:

1. Log out all the subscribers.
2. Disable the interface.
3. Make the desired attribute changes.
4. Reenable the interface.

If you do not follow these steps, the attribute change brings down the physical interface and all subscribers using that interface.

To avoid service interruptions, we recommend that you make the changes during a maintenance window.

RELATED DOCUMENTATION

Understanding Hierarchical CoS for Subscriber Interfaces

Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links

[Configuring an Interface Set of Subscribers in a Dynamic Profile | 192](#)

Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview

Static and Dynamic VLAN Subscriber Interfaces over Aggregated Ethernet Overview

Distribution of Demux Subscribers in an Aggregated Ethernet Interface

CoS for PPPoE Subscriber Interfaces Overview

For all supported hardware platforms, you can attach an output traffic-control profile that contains basic shaping and scheduling properties directly to a static or dynamic 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.

For MPCs that support hierarchical scheduling, you can shape subscriber or access node traffic at different levels of the PPPoE interface hierarchy by attaching 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.

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 demux 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 Hierarchical CoS for Subscriber Interfaces

Configuring Static Hierarchical Scheduling in a Dynamic Profile

Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface

CoS on Enhanced IQ2 PICs Overview

Configuring Scheduling and Shaping for Subscriber Access

IN THIS CHAPTER

- [Configuring Traffic Scheduling and Shaping for Subscriber Access | 47](#)
- [Configuring Schedulers in a Dynamic Profile for Subscriber Access | 52](#)
- [Configuring Scheduler and Scheduler Map Sharing | 60](#)
- [Example: Providing Unique Rate Configurations for Schedulers in a Dynamic Profile | 62](#)
- [Example: Configuring Aggregate Scheduling of Queues for Residential Subscribers on Static IP Demux Interfaces | 63](#)
- [Verifying the Scheduling and Shaping Configuration for Subscriber Access | 66](#)

Configuring Traffic Scheduling and Shaping for Subscriber Access

IN THIS SECTION

- [Configuring Static Traffic Shaping and Scheduling Parameters in a Dynamic Profile | 48](#)
- [Configuring Dynamic Traffic Shaping and Scheduling Parameters in a Dynamic Profile | 49](#)
- [Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers | 50](#)

You use traffic-control profiles to configure traffic shaping and scheduling properties.

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

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. Apply a static scheduler map that has been 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
```

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)
```

SEE ALSO

Configuring Static Hierarchical Scheduling in a Dynamic Profile

Example: Maintaining a Constant Traffic Flow by Configuring a Static VLAN Interface with a Dynamic Profile for Subscriber Access

Example: Configuring Dynamic Hierarchical Scheduling for Subscribers

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. Reference a dynamic scheduler map.

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 $junos-cos-scheduler-map
```

3. Configure the shaping rate variable.

```
[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
```

SEE ALSO

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

Example: Maintaining a Constant Traffic Flow by Configuring a Static VLAN Interface with a Dynamic Profile for Subscriber Access

Example: Configuring Dynamic Hierarchical Scheduling for Subscribers

[CoS for Subscriber Access Overview | 39](#)

Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers

IN THIS SECTION

- [Using the CLI to Globally Modify a Traffic-Control Profile Currently Applied to Multiple Subscribers | 51](#)
- [Using the CLI to Modify a Traffic-Control Profile for a Specific Current Subscriber | 51](#)

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. See *Junos OS Predefined Variables and Configuring Predefined Dynamic Variables in Dynamic Profiles* for more information about predefined variables.

There are two methods you can use to modify a 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.

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

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
```

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.

SEE ALSO

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

RELATED DOCUMENTATION

For hardware requirements and configuration guidelines, see [Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[CoS for Subscriber Access Overview | 39](#)

Configuring Static Hierarchical Scheduling in a Dynamic Profile

Example: Maintaining a Constant Traffic Flow by Configuring a Static VLAN Interface with a Dynamic Profile for Subscriber Access

Example: Configuring Dynamic Hierarchical Scheduling for Subscribers

[Verifying the Scheduling and Shaping Configuration for Subscriber Access | 66](#)

Configuring Schedulers in a Dynamic Profile for Subscriber Access

IN THIS SECTION

- [Configuring Static Schedulers in a Dynamic Profile | 53](#)
- [Configuring Dynamic Schedulers with Variables in a Dynamic Profile | 54](#)
- [Configuring a Combination of Static and Dynamic Scheduler Parameters in a Scheduler Definition | 56](#)

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

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
```

- 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
```

SEE ALSO

For hardware requirements and configuration guidelines, see [Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[Changing CoS Services Overview | 160](#)

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
```

SEE ALSO

For hardware requirements and configuration guidelines, see [Guidelines for Configuring Dynamic CoS for Subscriber Access](#) | 40

[Changing CoS Services Overview](#) | 160

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]
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
```

SEE ALSO

For hardware requirements and configuration guidelines, see [Guidelines for Configuring Dynamic CoS for Subscriber Access](#) | 40

RELATED DOCUMENTATION

For hardware requirements and configuration guidelines, see [Guidelines for Configuring Dynamic CoS for Subscriber Access](#) | 40

[Verifying the Scheduling and Shaping Configuration for Subscriber Access](#) | 66

[Changing CoS Services Overview](#) | 160

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 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*.

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.

```
[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

Dynamic Profiles Overview

[Guidelines for Configuring Dynamic CoS for Subscriber Access](#) | 40

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 {
    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](#) | 56

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 multiplex 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 192.0.2.4/24;
      }
    }
  }
}
demux0 {
  unit 0 {
    demux-options {
      underlying-interface ge-2/0/1.1;
    }
    family inet {
      address 192.0.2.1/24;
      demux-source {
        192.0.2.0/24;
      }
    }
  }
  unit 1 {
    demux-options {
      underlying-interface ge-2/0/1.1;
    }
    family inet {
      address 192.0.2.21/24;
      demux-source {
        192.0.2.20/24;
      }
    }
  }
}
```

```

    }
}
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 {
      transmit-rate percent 30;

```

```

        buffer-size percent 30;
    }
    s3 {
        transmit-rate percent 40;
        buffer-size percent 40;
    }
}
}

```

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
```

- 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
```

Configuring Hierarchical CoS Scheduling on MPLS Ethernet Pseudowire Subscriber Interfaces

IN THIS CHAPTER

- [Enhanced Subscriber Management Subscriber Logical Interfaces or Interface Sets Over Underlying Logical Interfaces for a CoS scheduler Hierarchy | 67](#)
- [Enhanced Subscriber Management Subscriber Logical Interfaces or Interface Sets Over MPLS Pseudowires for a CoS scheduler Hierarchy | 70](#)
- [Configuring Layer 2 Subscriber Logical Interfaces for CoS Hierarchical Schedulers Using Dynamic Profiles for Differentiating Home and Access Node Networks | 73](#)
- [Example: Configuring Layer 2 Subscriber Logical Interfaces for CoS Hierarchical Schedulers Using Static CoS for Differentiating Home and Access Node Networks | 79](#)

Enhanced Subscriber Management Subscriber Logical Interfaces or Interface Sets Over Underlying Logical Interfaces for a CoS scheduler Hierarchy

Starting in Junos OS Release 15.1, you can enable a CoS scheduling hierarchy for subscriber logical interfaces or interface sets over underlying logical interfaces. Until Junos OS Release 14.2, an interface set can be either at Layer 2 or Layer 3 levels of the CoS three-level hierarchical scheduler. When the interface set is at the Layer 3 level, a mechanism to configure the Layer 2 node to which the Layer 3 node belonged was not available. As a result, the Layer 2 node was a dummy node in such a case for the three-level hierarchical scheduler.

In certain Broadband Remote Access Server (B-RAS) deployments, when you use an interface set to denote a home network, it might be necessary to configure the home network and the access node (such as a digital subscriber line access multiplexer, or DSLAM) in a scheduler hierarchy. This method of hierarchical scheduler is necessary in agent circuit identifier (ACI) VLANs because a home or an ACI is always an interface set in such topologies. You can now enable an enhanced subscriber management logical interface, such as an MPLS pseudowire logical interface to function as a Layer 2 node in a CoS hierarchical scheduler. A subscriber logical interface is considered to operate at Layer 2 only if you configure three-level hierarchical scheduling on the logical tunnel anchor point on the physical interface (the IFD). An MPLS pseudowire is a virtual device that is stacked above the logical tunnel anchor point. Implicit hierarchy processes the interface stack properly in such a setup. To configure three-level hierarchical scheduling, include the

implicit-hierarchy option at the `[edit interfaces "$junos-interface-ifd-name" hierarchical-scheduler]` or the `[edit interfaces lt-device hierarchical-scheduler]` hierarchy level. 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. Subscriber logical interfaces at Layer 3 that are stacked over the underlying logical interfaces at Layer 2 are supported if the Layer 2 logical interface is an underlying interface of the Layer 3 interface.

For example, if a PPPoE logical interface contains an underlying logical interface, ge-1/0/0.100, the ge-1/0/0.100 interface can be at Layer 2 and the PPPoE logical interface can be at Layer 3. You can also configure PPP or IP demux interfaces in such a fashion at Layer 3. Similarly, you can configure logical interfaces at Layer 2 that serve as underlying interfaces for logical interface sets, such as PPPoE ACI interface sets or IP demux interface sets, where all the member logical interfaces of the interface set contain the same underlying logical interface at Layer 2. You can configure the logical interfaces at Layer 2 in a dynamic profile or in a static CoS configuration.

Dynamic profile CoS configuration for underlying logical interfaces is supported because two interface stanzas with TCPs in one dynamic profile are considered valid. For dynamic underlying logical interfaces, you can configure in a profile different from the client logical interface profile or in the same profile as the client profile. If the underlying logical interface is static and CoS is configured dynamically in a dynamic profile, it must be specified in the same profile as the client logical interface. However, CoS for the underlying logical interfaces must be configured either in a dynamic profile or in a static CoS because both static CoS and dynamic CoS are not supported on the same logical interface.

Reparenting is a technique that denotes the movement of the CoS hierarchical scheduler from one node to another node, such as moving all logical interfaces stacked over an underlying logical interface on top of the base physical interface to be over the underlying logical interface directly and adding the scheduling node. This movement might occur when when CoS for the underlying logical interface or the underlying interface set is configured later than the client logical interface (IP demux or PPPoE).

Reparenting is not supported for enhanced subscriber management logical interfaces in a CoS hierarchical scheduler that includes enhanced subscriber management logical interfaces over a purely dynamic column and enhanced subscriber management logical interfaces over a partially static column. The following describes real-world network environments where reparenting might be required and the alternative approaches that can be adopted in such conditions:

Adding or removing static CoS configuration from an IFL set or an underlying IFL with enhanced subscriber management logical interface on top of it—In such a scenario, adding or removing static CoS is not supported after a subscriber has logged in to the interface column in an environment where enhanced subscriber management is enabled. A commit error occurs when you attempt this CoS configuration change. This

commit failure is not a problem in customer networks because the networks are previously designed, Layer 2 nodes specified, and CoS is configured much before clients are logged in.

Two dynamic profiles for Client logical interfaces over a single CVLAN (or an ACI VLAN) with underlying CoS configuration in one client profile and not in the other profile—In such a scenario, you can maintain dynamic profiles with underlying configuration to be consistent – either all profiles contain underlying CoS config or none of them contain CoS configuration. A negative acknowledgment is sent when a subscriber attempts to log in if a differing way of CoS configuration is observed in the client profiles.

A client profile for an internal node (for example, C-VLAN or IFL set) that does not contain CoS initially and CoS is applied later using a service profile—In such a scenario, it is required that you always specify CoS scheduling in the client profile if you want to reapply some of the settings using a service profile. If this method of configuration is not adopted, a negative acknowledgment is sent when a subscriber attempts to log in. Static or dynamic demux, PPPoE, or PPP interfaces over aggregated Ethernet logical interfaces are not supported.

Consider a scenario in which three subscriber queues, namely, PPPoE subscriber queue 1, PPPoE subscriber queue 2, and DHCP subscriber queues, are established. A Gigabit Ethernet interface, ge-1/0/0 is at Layer 1. Two Layer 2 interface nodes are stacked over the Layer 1 base interface. The Layer 2 interfaces are ge-1/0/0.x or demux0.x and ge-1/0/0.y or demux0.y. Logical interface sets, pppoe-iflset (for access node) and demux-iflset (for home network), are configured at Layer 3 to handle two sets of PPPoE subscriber queues respectively over the Layer 2 interface, ge-1/0/0.x or demux0.x. A traffic control profile, subscriber-tcp, is attached to both these Layer 3 IFL sets. ppp-demux-iflset (demux and pppoe) is the interface set over the Layer 2 interface of ge-1/0/0.y or demux0.y. A traffic control profile, subscriber-tcp, is attached to this interface set. ge-1/0/0.X or demux0.X is the UIFL for all logical interfaces that belong to the pppoe-iflset and demux-iflset. In this topology, ge-1/0/0.Y or demux0.Y is the UIFL for all logical interface that belong to ppp-demux-iflset.

Consider another scenario in which three subscriber queues, PPPoE subscriber queues, demux subscriber queues, and DHCP subscriber queues, are established. A Gigabit Ethernet interface, ge-1/0/0 is at Layer 1. Two Layer 2 interface nodes are stacked over the Layer 1 base interface. The Layer 2 interfaces are ge-1/0/0.X or demux0.X, and ge-1/0/0.Y or demux0.Y. At Layer 3, pp0.XX is configured over the underlying Layer 2 interface of ge-1/0/0.X or demux0.X, demux0.ZZ is configured over the underlying Layer 2 interface of ge-1/0/0.X or demux0.X, and pp0.YY is configured over the underlying Layer 2 interface of ge-1/0/0.Y or demux0.Y. Traffic control profiles, subscriber-tcp, are applied to pp0.xx for PPPoE subscriber queues, to demux0.yy for demux subscriber queues, and pp0.yy for DHCP subscriber queues. In this topology, ge-1/0/0.X or demux0.X is the underlying IFL for pp0.XX and demux0.ZZ. ge-1/0/0.Y or demux0.Y is the underlying IFL for pp0.YY.

RELATED DOCUMENTATION


Enhanced Subscriber Management Subscriber Logical Interfaces or Interface Sets Over MPLS Pseudowires for a CoS scheduler Hierarchy

Starting in Junos OS Release 15.1, you can enable a CoS scheduling hierarchy for subscriber logical interfaces or interface sets over underlying MPLS pseudowire logical interfaces. Until Junos OS Release 14.2, an interface set can be either at Layer 2 or Layer 3 levels of the CoS three-level hierarchical scheduler. When the interface set is at the Layer 3 level, a mechanism to configure the Layer 2 node to which the Layer 3 node belonged was not available. As a result, the Layer 2 node was a dummy node in such a case for the three-level hierarchical scheduler.

In certain Broadband Remote Access Server (B-RAS) deployments, when you use an interface set to denote a home network, it might be necessary to configure the home network and the access node (such as a digital subscriber line access multiplexer, or DSLAM) in a scheduler hierarchy. This method of hierarchical scheduler is necessary in agent circuit identifier (ACI) VLANs because a home or an ACI is always an interface set in such topologies.

Enhanced subscriber management enables you to take advantage of increased scaling and performance for configuring and managing dynamic interfaces and services for subscriber management. You can now enable an enhanced subscriber management logical interface, such as an MPLS pseudowire logical interface to function as a Layer 2 node in a CoS hierarchical scheduler. A subscriber logical interface is considered to operate at Layer 2 only if you configure three-level hierarchical scheduling on the logical tunnel anchor point on the physical interface (the IFD). An MPLS pseudowire is a virtual device that is stacked above the logical tunnel anchor point. Implicit hierarchy processes the interface stack properly in such a setup. To configure three-level hierarchical scheduling, include the **implicit-hierarchy** option at the **[edit interfaces "\$junos-interface-afd-name" hierarchical-scheduler]** or the **[edit interfaces lt-device hierarchical-scheduler]** hierarchy level. 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. [Figure 1 on page 71](#) shows the protocol stack for a pseudowire subscriber logical interface.

Figure 1: Pseudowire Subscriber Interface Protocol Stack



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. Subscriber logical interfaces at Layer 3 that are stacked over the underlying MPLS pseudowire logical interfaces at Layer 2 are supported if the Layer 2 logical interface is an underlying interface of the Layer 3 interface.

For example, if a PPPoE logical interface contains an MPLS pseudowire, psp-anch-device-name.logical-unit-number, as the underlying interface, the psp-anch-device-name.logical-unit-number interface can be at Layer 2 and the PPPoE logical interface can be at Layer 3. You can also configure PPP or IP demux interfaces in such a fashion at Layer 3. Similarly, you can configure MPLS pseudowire logical interfaces at Layer 2 that serve as underlying interfaces for logical interface sets, such as PPPoE ACI interface sets or IP demux interface sets, where all the member logical interfaces of the interface set contain the same underlying MPLS pseudowire at Layer 2. You can configure the MPLS pseudowire logical interfaces at Layer 2 in a dynamic profile or in a static CoS configuration.

Dynamic profile CoS configuration for underlying logical interfaces is supported because two interface stanzas with TCPs in one dynamic profile are considered valid. For dynamic pseudowire underlying logical interfaces, you can configure in a profile different from the client logical interface profile or in the same profile as the client profile. If the underlying logical interface is static and CoS is configured dynamically in a dynamic profile, it must be specified in the same profile as the client logical interface. However, CoS

for the underlying logical interfaces must be configured either in a dynamic profile or in a static CoS because both static CoS and dynamic CoS are not supported on the same logical interface.

Reparenting is a technique that denotes the movement of the CoS hierarchical scheduler from one node to another node, such as moving all logical interfaces stacked over an underlying logical interface on top of the base physical interface to be over the underlying logical interface directly and adding the scheduling node. This movement might occur when when CoS for the underlying logical interface or the underlying interface set is configured later than the client logical interface (IP demux or PPPoE).

Reparenting is not supported for enhanced subscriber management logical interfaces in a CoS hierarchical scheduler that includes enhanced subscriber management logical interfaces over a purely dynamic column and enhanced subscriber management logical interfaces over a partially static column. The following describes real-world network environments where reparenting might be required and the alternative approaches that can be adopted in such conditions:

Adding or removing static CoS configuration from a logical interface (IFL) set or an underlying IFL with enhanced subscriber management logical interface on top of it is not supported. In such a scenario, adding or removing static CoS is not supported after a subscriber has logged in to the interface column in an environment where enhanced subscriber management is enabled. A commit error occurs when you attempt this CoS configuration change. This commit failure is not a problem in customer networks because the networks are previously designed, Layer 2 nodes specified, and CoS is configured much before clients are logged in.

Two dynamic profiles for Client logical interfaces over a single CVLAN (or an ACI VLAN) with underlying CoS configuration in one client profile and not in the other profile—In such a scenario, you can maintain dynamic profiles with underlying configuration to be consistent – either all profiles contain underlying CoS config or none of them contain CoS configuration. A negative acknowledgment is sent when a subscriber attempts to log in if a differing way of CoS configuration is observed in the client profiles.

A client profile for an internal node (for example, C-VLAN or IFL set) that does not contain CoS initially and CoS is applied later using a service profile—In such a scenario, it is required that you always specify CoS scheduling in the client profile if you want to reapply some of the settings using a service profile. If this method of configuration is not adopted, a negative acknowledgment is sent when a subscriber attempts to log in. Static or dynamic demux, PPPoE, or PPP interfaces over aggregated Ethernet logical interfaces are not supported.

Consider a scenario in which three subscriber queues, namely, PPPoE subscriber queue 1, PPPoE subscriber queue 2, and DHCP subscriber queues, are established. A logical tunnel interface, It-1/0/0 is at Layer 1. Two Layer 2 interface nodes are stacked over the Layer 1 base interface. The Layer 2 interfaces are psX.Y and psX.Z. Logical interface sets, ppp0.XX (for access node) and demux0.ZZ (for home network), are configured at Layer 3 to handle PPPoE subscriber queues and DHCP subscriber queues respectively over the Layer 2 interface, psX.Y. A logical interface, pp0.YY, is configured at Layer 3 to handle PPPoE subscriber queues over the Layer 2 interface, psX.Z. A traffic control profile, subscriber-tcp, is attached to these Layer 3 interfaces. psX.Y is the underlying logical interface for pp0.XX and demux0.ZZ if Y is not 0. psX.Z is the

underlying logical interface for pp0.YY if Z is not 0. psX.0 is called the pseudowire transport logical interface and psX.Y (where Y is not equal to 0) is called the pseudowire service logical interface.

Consider another scenario in which two subscriber queues, PPPoE subscriber queues and DHCP subscriber queues, are established. A logical tunnel interface, lt- 1/0/0 is at Layer 1. Two Layer 2 interface nodes are stacked over the Layer 1 base interface. The Layer 2 interfaces are psX.Y and psX.Z. Logical interface sets, pppoe-iflset (for access node) and demux-iflset (for home network), are configured at Layer 3 to handle PPPoE subscriber queues and DHCP subscriber queues respectively over the Layer 2 interface, psX.Y. A logical interface set, ppp-demux-iflset, is configured at Layer 3 to handle PPPoE and DHCP subscriber queues over the Layer 2 interface, psX.Z. A traffic control profile, subscriber-tcp, is attached to these Layer 3 interfaces. psX.Y is the underlying logical interface for all logical interfaces that belong to the pppoe-iflset and demux-iflset if Y is not equal to 0. psX.Z is the underlying logical interface for all logical interfaces that belong to the ppp-demux-iflset interface set if Z is not 0. psX.0 is called the pseudowire transport logical interface and psX.Y (where Y is not equal to 0) is called the pseudowire service logical interface.

RELATED DOCUMENTATION

Configuring Layer 2 Subscriber Logical Interfaces for CoS Hierarchical Schedulers Using Dynamic Profiles for Differentiating Home and Access Node Networks

In certain Broadband Remote Access Server (B-RAS) deployments, when you use an interface set to denote a home network, it might be necessary to configure the home network and the access node (such as a digital subscriber line access multiplexer, or DSLAM) in a scheduler hierarchy. This method of hierarchical scheduler is necessary in agent circuit identifier (ACI) VLANs because a home or an ACI is always an interface set in such topologies. You can configure a subscriber logical interface or an interface set at Layer 3 over an underlying enhanced subscriber management logical interface that functions as a Layer 2 node. You can configure a the Layer 2 logical interface in a CoS dynamic profile.

Consider a scenario in which a Layer 3 interface set, ACI-set aci-1006-ps0.3221225479, is stacked over dynamic a MPLS pseudowire service logical interface, ps0.3221225479, at Layer 2. 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. Two traffic control profiles are defined to apply an output traffic scheduling and shaping profile to the MPLS pseudowire logical interface. These control profiles are an-tcp to be applied for TCP subscribers that are terminated at the access mode and an-tcp-remaining, which is 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.

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*.

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
user@host# set shaping-rate rate
user@host# set guaranteed-rate rate
```

4. Specify the output traffic control profile and the remaining traffic control profile for the underlying logical interfaces that are members of the interface set.

```
[edit class-of-service interfaces]
user@host# edit interface "$junos-interface-ifd-name" unit "$junos-underlying-interface-unit"
user@host# edit output-traffic-control-profile profile-name
user@host# edit output-traffic-control-profile-remaining profile-name
```

5. Specify the output traffic control profile for the interface set.

```
[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;
    }
}

```

```

    }
}

```

You can use the **show class-of-service scheduler-hierarchy interface *interface-name*** command to verify the CoS hierarchical schedulers configured on the interfaces. For example, the following output illustrates that ACI-set aci-1003-demux0.3221225482 is stacked over demux0.3221225482.

```

user@host> show class-of-service scheduler-hierarchy interface ge-0/2/0
Interface/
Resource name      Shaping Guaranteed  Guaranteed/  Queue  Excess
                   rate    rate             Excess      weight weight
                   kbits   kbits           priority   high/low
ge-0/2/0           1000000
  ge-0/2/0 RTP      1000000           0
    best-effort      1000000           0      Low  Low    950
    network-control   1000000           0      Low  Low    50
  demux0.3221225482  100000           80000
    demux0.3221225482 RTP
      30000           20000
        best-effort    30000           19000      Low  Low    950
        network-control 30000           1000      Low  Low    50
  aci-1003-demux0.3221225482 out-of-scheduler-resources

```

From the following sample output, you can verify that ACI-iflset aci-1001-ps1.3221225472 is stacked over a static pseudowire transport logical interface, ps1.0

```

user@host> show class-of-service scheduler-hierarchy interface ps1
Interface/
Resource name      Shaping Guaranteed  Guaranteed/  Queue  Excess
                   rate    rate             Excess      weight weight
                   kbits   kbits           priority   high/low
lt-0/3/0           10000000
  lt-0/3/0 RTP      10000000           0
    best-effort      10000000           0      Low  Low    950
    network-control   10000000           0      Low  Low    50
  ps1.0             100000           0
    ps1.0 RTP        500000           0
      best-effort      400000           0      Low  Low    1000
  aci-1001-ps1.3221225472
    200000           10000
      best-effort      160000           2000      Low  Low    1000

```

From the following sample output, you can verify that ACI-set aci-1006-ps0.3221225479 is stacked over the dynamic pseudowire service logical interface, ps0.3221225479.

```

user@host> show class-of-service scheduler-hierarchy interface ps0

```

Interface/ Resource name	Shaping rate kbits	Guaranteed rate kbits	Guaranteed/ Excess priority	Queue weight	Excess weight high/low
lt-0/3/0	10000000				
lt-0/3/0 RTP	10000000	0			1 1
best-effort	10000000	0	Low Low	950	
network-control	10000000	0	Low Low	50	
ps0.32767	10000000	2000			50 50
best-effort	10000000	1900	Low Low	950	
network-control	10000000	100	Low Low	50	
ps0.3221225479	100000	0			1 1
ps0.3221225479 RTP	40000	20000			500 500
best-effort	5000	3000	Medium Low	1	
expedited-forwarding	40000	2000	Medium High	1000	
aci-1006-ps0.3221225479					
	100000	10000			250 250
best-effort	5000	1500	Medium Low	1	
expedited-forwarding	100000	1000	Medium High	500	
assured-forwarding	100000	1000	Medium High	500	
network-control	100000	2000	High High	1	

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[Changing CoS Services Overview | 160](#)

Example: Configuring Layer 2 Subscriber Logical Interfaces for CoS Hierarchical Schedulers Using Static CoS for Differentiating Home and Access Node Networks

IN THIS SECTION

- [Requirements | 79](#)
- [Overview | 79](#)
- [Configuration | 80](#)
- [Verification | 82](#)

Starting in Junos OS Release 15.1, in certain Broadband Remote Access Server (B-RAS) deployments, when you use an interface set to denote a home network, it might be necessary to configure the home network and the access node (such as a digital subscriber line access multiplexer, or DSLAM) in a scheduler hierarchy. This method of hierarchical scheduler is necessary in agent circuit identifier (ACI) VLANs because a home or an ACI is always an interface set in such topologies. You can enable an enhanced subscriber management logical interface, such as an MPLS pseudowire logical interface to function as a Layer 2 node in a CoS hierarchical scheduler. A subscriber logical interface is considered to operate at Layer 2 only if you configure CoS three-level hierarchical scheduling on the logical tunnel anchor point on the physical interface (the IFD). 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.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 15.1
- MX Series Router with MPCs

Overview

You specify an anchor point, which identifies the logical tunnel interface that terminates the pseudowire tunnel at the access node. Consider a scenario in which `lt-0/3/0` is the logical tunnel interface, and an MPLS pseudowire transport logical interface, `ps1.0`, that is anchored on the logical tunnel. Three-level hierarchical scheduling is enabled on the logical tunnel interface for static CoS configuration.

Configuration

IN THIS SECTION

- [Configuring an MPLS Pseudowire Transport Logical Interface Over a Logical Tunnel in a Static CoS Setup | 80](#)
- [Results | 82](#)

To configure an enhanced subscriber management logical interface, such as an MPLS pseudowire logical interface to function as a Layer 2 node in a CoS hierarchical scheduler, perform these tasks:

CLI Quick Configuration

To quickly configure the MPLS pseudowire logical interface to function as a Layer 2 node in a three-level hierarchical scheduler, copy the following commands and paste them into the router terminal window:

```
[edit]
set interfaces lt-0/3/0
set interfaces lt-0/3/0 hierarchical-scheduler implicit-hierarchy
set interfaces ps1
set interfaces ps1 description client-port-l2circuit
set interfaces ps1 flexible-vlan-tagging
set interfaces ps1 anchor-point lt-0/3/0
set interfaces ps1 unit 0
set interfaces ps1 unit 0 encapsulation ethernet-ccc
set interfaces ps1 unit 0 output-traffic-control-profile profile-name
```

Configuring an MPLS Pseudowire Transport Logical Interface Over a Logical Tunnel in a Static CoS Setup

Step-by-Step Procedure

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 three-level scheduling on pseudowire transport logical interfaces over a logical tunnel physical interface (LT ifd):

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 lt-0/3/0
user@host#set hierarchical-scheduler implicit-hierarchy
```

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

```
[edit]
user@host# edit interfaces ps1
```

3. Configure a description for the pseudowire subscriber logical interface.

```
[edit interfaces ps1]
user@host# set description client-port-l2circuit
```

4. Specify the **flexible-vlan-tagging** statement to indicate that this interface is for use with both VLAN and stacked VLAN ranges.

```
[edit interfaces ps1]
user@host# set flexible-vlan-tagging
```

5. Specify the logical tunnel (lt) interface that identifies the Packet Forwarding Engine that processes the pseudowire termination.

```
[edit interfaces ps1]
user@host# set anchor-point lt-0/3/0
```

6. Specify that you want to configure unit 0, which represents the transport logical interface.

```
[edit interfaces ps1]
user@host# edit unit 0
```

7. Specify the ethernet-ccc encapsulation method for the transport logical interface.

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

8. Specify the traffic-control profile to use on the pseudowire transport logical interface.

```
[edit class-of-service]
user@host#edit interfaces ps 1
user@host#edit unit 0
user@host#set output-traffic-control-profile profile-name
```

Results

In configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lt-0/3/0 {
    hierarchical-scheduler implicit-hierarchy;
}

ps1 {
    description client-port-l2circuit;
    anchor-point {
        lt-0/3/0;
    }
    flexible-vlan-tagging;
    unit 0 {
        encapsulation ethernet-ccc;
    }
}
```

Verification

IN THIS SECTION

- [Verifying the Scheduler Hierarchy Configured on the Interfaces | 82](#)

Confirm that the configuration is working properly.

Verifying the Scheduler Hierarchy Configured on the Interfaces

Purpose

Verify the CoS hierarchical scheduler configured on the Layer 2 and Layer 3 interface nodes.

Action

From operational mode, enter the **show class-of-service scheduler-hierarchy interface ps0** command.

```
user@host> show class-of-service scheduler-hierarchy interface ps0
```

Interface/ Resource name	Shaping rate kbits	Guaranteed rate kbits	Guaranteed/ Excess priority		Queue weight	Excess weight high/low	
lt-0/3/0	10000000						
lt-0/3/0 RTP	10000000	0				1	1
best-effort	10000000	0	Low	Low	950		
network-control	10000000	0	Low	Low	50		
ps0.0	200000	0				1	1
ps0.0 RTP	10000000	0				1	1
best-effort	10000000	0	Low	Low	950		
network-control	10000000	0	Low	Low	50		
ps0.3221225474	100000	0				1	1
best-effort	5000	0	Medium	Low	1000		

user@host> show class-of-service scheduler-hierarchy interface ps0

Interface/ Resource name	Shaping rate kbits	Guaranteed rate kbits	Guaranteed/ Excess priority		Queue weight	Excess weight high/low	
lt-0/3/0	10000000						
lt-0/3/0 RTP	10000000	0				1	1
best-effort	10000000	0	Low	Low	950		
network-control	10000000	0	Low	Low	50		
ps0.32767	10000000	2000				33	33
best-effort	10000000	1900	Low	Low	950		
network-control	10000000	100	Low	Low	50		
ps0.3221225474	200000	0				1	1
ps0.3221225474 RTP	100000	30000				500	500
best-effort	30000	3000	Medium	Low	250		
expedited-forwarding	32000	9000	Low	Low	750		
pp0.3221225475	100000	10000				166	166
best-effort	5000	1500	Low	Low	1		
expedited-forwarding	100000	1000	Medium	High	500		
assured-forwarding	100000	1000	Medium	High	500		
network-control	100000	2000	High	High	1		

Meaning

Shows that dynamic pseudowire service logical interface, ps0.3221225474, is stacked over the static pseudowire transport logical interface, ps0.0. Also, the sample output denotes that pp0.3221225475 is stacked over dynamic pseudowire service logical interface, ps0.3221225474.

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, in certain Broadband Remote Access Server (B-RAS) deployments, when you use an interface set to denote a home network, it might be necessary to configure the home network and the access node (such as a digital subscriber line access multiplexer, or DSLAM) in a scheduler hierarchy.

RELATED DOCUMENTATION

Allocating Dedicated Queues for Each Logical Interface Using Per-Unit Scheduling

IN THIS CHAPTER

- [Hardware Requirements for Dynamic Per-Unit Scheduling | 85](#)
- [Configuring Per-Unit Scheduling in a Dynamic Profile | 86](#)
- [Example: Configuring Per-Unit Scheduling for Subscriber Access | 88](#)

Hardware Requirements for Dynamic Per-Unit Scheduling

Table 12 on page 85 lists the hardware requirements based on subscriber interface type for per-unit scheduling in dynamic CoS configurations.

Table 12: Hardware Required for Per-Unit Scheduling Dynamic CoS Configurations

Subscriber Interface Type	EQ DPCs on MX Series Routers	MPC/MIC Modules on MX Series Routers
Static and dynamic VLANs	Yes	Yes
Static and dynamic VLANs over aggregated Ethernet	No	No
Static or dynamic IP demux interfaces	Yes	No
Static or dynamic IP demux interfaces over aggregated Ethernet	No	No
Static or dynamic VLAN demux interfaces	No	No
Static or dynamic VLAN demux interfaces over aggregated Ethernet	No	No

Table 12: Hardware Required for Per-Unit Scheduling Dynamic CoS Configurations (*continued*)

Subscriber Interface Type	EQ DPCs on MX Series Routers	MPC/MIC Modules on MX Series Routers
Static PPPoE interfaces	No	Yes
Dynamic PPPoE interfaces	No	No
Static or dynamic PPPoE interfaces over aggregated Ethernet	No	No
L2TP LAC tunnel over PPP	No	No
L2TP LNS inline service over PPP	No	No

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[Configuring Per-Unit Scheduling in a Dynamic Profile | 86](#)

Configuring Per-Unit Scheduling in a Dynamic Profile

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 *Defining Packet Drop Behavior by Configuring RED Drop Profiles*.

- c. Configure the forwarding classes.

See *Configuring a Custom Forwarding Class for Each Queue*.

- d. Configure the rewrite-rules and classifier definitions.

See *Configuring Rewrite Rules* and *Configuring Behavior Aggregate Classifiers*.

See *The Junos OS CoS Components Used to Manage Congestion and Control Service Levels* for information about configuring the remaining CoS parameters.

2. Configure a static or dynamic subscriber interface that can be referenced in the dynamic profile.

3. Configure CoS parameters in a dynamic profile.

- a. Configure the dynamic profile.

See *Configuring a Basic Dynamic Profile*.

- 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 47](#).

- 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 52](#).

- 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 217](#).
- For rewrite rules, see [“Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile” on page 219](#).
- For classifiers, see [“Applying a Classifier to a Subscriber Interface in a Dynamic Profile” on page 221](#).

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 Static Default Values for Traffic Scheduling and Shaping” on page 168](#)

- 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*.

RELATED DOCUMENTATION

[CoS for Subscriber Access Overview | 39](#)

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[Example: Configuring Per-Unit Scheduling for Subscriber Access | 88](#)

Example: Configuring Per-Unit Scheduling for Subscriber Access

In this example, a network administrator sets up a subscriber access configuration with per-unit scheduling.

1. The administrator configures the static VLAN interfaces and enables per-unit scheduling for the interfaces.

```
[edit]
interfaces {
  ge-1/1/0 {
    per-unit-scheduler;
    vlan-tagging;
```

```

unit 100 {
    vlan-id 100;
    family inet {
        unnumbered-address lo0.0 preferred-source-address 192.0.2.100;
    }
}
unit 200 {
    vlan-id 200;
    family inet {
        unnumbered-address lo0.0 preferred-source-address 192.0.2.100;
    }
}
}
ge-1/1/1 {
    per-unit-scheduler;
    vlan-tagging;
    unit 100 {
        vlan-id 100;
        family inet {
            unnumbered-address lo0.0 preferred-source-address 192.0.2.100;
        }
    }
    unit 200 {
        vlan-id 200;
        family inet {
            unnumbered-address lo0.0 preferred-source-address 192.0.2.100;
        }
    }
}
ge-1/0/1 {
    unit 0 {
        family inet {
            address 203.0.113.31/24;
        }
    }
}
ge-1/1/2 {
    description "wfce14 eth1 soso ge-1/1/2";
    vlan-tagging;
    gigether-options {
        no-auto-negotiation;
    }
    unit 100 {
        vlan-id 100;
    }
}

```

```

    family inet {
        address 203.0.113.121/24;
    }
}
}
}

```

2. The administrator configures static CoS parameters, including forwarding classes and classifiers, to be referenced in the dynamic profiles.

```

[edit]
class-of-service {
  classifiers {
    inet-precedence 8q-inet {
      forwarding-class be {
        loss-priority low code-points 000;
      }
      forwarding-class ef {
        loss-priority low code-points 001;
      }
      forwarding-class af {
        loss-priority low code-points 010;
      }
      forwarding-class nc {
        loss-priority low code-points 011;
      }
      forwarding-class voice {
        loss-priority low code-points 100;
      }
      forwarding-class video {
        loss-priority low code-points 101;
      }
      forwarding-class game {
        loss-priority low code-points 110;
      }
      forwarding-class data {
        loss-priority low code-points 111;
      }
    }
    inet-precedence 4q-inet {
      forwarding-class be {
        loss-priority low code-points [ 000 001 ];
      }
    }
  }
}

```

```

forwarding-class ef {
    loss-priority low code-points [ 010 011 ];
}
forwarding-class af {
    loss-priority low code-points [ 100 101 ];
}
forwarding-class nc {
    loss-priority low code-points [ 110 111 ];
}
}
inet-precedence 8q-drop-inet {
    forwarding-class be {
        loss-priority low code-points 000;
    }
    forwarding-class ef {
        loss-priority medium-low code-points 001;
    }
    forwarding-class af {
        loss-priority medium-high code-points 010;
    }
    forwarding-class nc {
        loss-priority high code-points 011;
    }
    forwarding-class voice {
        loss-priority low code-points 100;
    }
    forwarding-class video {
        loss-priority medium-low code-points 101;
    }
    forwarding-class game {
        loss-priority medium-high code-points 110;
    }
    forwarding-class data {
        loss-priority high code-points 111;
    }
}
inet-precedence 4q-drop-inet {
    forwarding-class be {
        loss-priority low code-points [ 000 001 ];
    }
    forwarding-class ef {
        loss-priority medium-low code-points [ 010 011 ];
    }
    forwarding-class af {

```

```

        loss-priority medium-high code-points [ 100 101 ];
    }
    forwarding-class nc {
        loss-priority high code-points [ 110 111 ];
    }
}
}
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 100 drop-probability 100;
        fill-level 0 drop-probability 0;
    }
    all {
        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;
    queue 4 voice;
    queue 5 video;
    queue 6 game;
    queue 7 data;
}
interfaces {
    ge-1/0/1 {
        unit 0 {
            classifiers {
                inet-precedence 8q-drop-low-high-inet;
            }
        }
    }
}

```

```

    }
  }
}
}
traceoptions {
  flag all;
  flag asynch;
  flag route-socket;
}
}

```

3. The administrator configures the access and service dynamic profiles to receive CoS parameters for the subscriber interfaces through RADIUS.

```

[edit]
dynamic-profiles {
  subscriber {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit" {
          family inet;
        }
      }
    }
  }
  class-of-service {
    traffic-control-profiles {
      zero {
        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" {
          output-traffic-control-profile zero;
        }
      }
    }
  }
  scheduler-maps {
    be_smap {
      forwarding-class be scheduler be_sch;
    }
  }
}

```

```

    }
    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 video scheduler video_sch;
        forwarding-class data scheduler data_sch;
    }
    be_ef_smap {
        forwarding-class be scheduler be_sch;
        forwarding-class ef scheduler ef_sch;
    }
    af_smap {
        forwarding-class af scheduler af_sch;
    }
    be_ef_af_nc_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;
    }
    voice_video_game_data_smap {
        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 {
    "$junos-cos-scheduler" {
        transmit-rate percent "$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 "$junos-cos-scheduler-dropfile-low";
        drop-profile-map loss-priority medium-low protocol any drop-profile
            "$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";
    }
}
}
}

```

```

service {
  variables {
    fc_1 default-value be;
    sch_1 default-value be_sch;
    sch-tx_1 default-value 20000000;
    sch-bs_1 default-value 10;
    sch-pri_1 default-value high;
    sch-drop-low_1 default-value d3;
    sch-drop-med-low_1 default-value d2;
    sch-drop-med-high_1 default-value d1;
    sch-drop-high_1 default-value d0;
    sch-drop-any_1 default-value d3;
    fc_2 default-value af;
    sch_2 default-value af_sch;
    sch-tx_2 default-value 10;
    sch-bs_2 default-value 10;
    sch-pri_2 default-value high;
    sch-drop-low_2 default-value d3;
    sch-drop-med-low_2 default-value d2;
    sch-drop-med-high_2 default-value d1;
    sch-drop-high_2 default-value d0;
    sch-drop-any_2 default-value d3;
    fc_3 default-value voice;
    sch_3 default-value voice_sch;
    sch-tx_3 default-value 20000000;
    sch-bs_3 default-value 10;
    sch-pri_3 default-value high;
    sch-drop-low_3 default-value d3;
    sch-drop-med-low_3 default-value d2;
    sch-drop-med-high_3 default-value d1;
    sch-drop-high_3 default-value d0;
    sch-drop-any_3 default-value d3;
    fc_4 default-value game;
    sch_4 default-value game_sch;
    sch-tx_4 default-value 10;
    sch-bs_4 default-value 10;
    sch-pri_4 default-value high;
    sch-drop-low_4 default-value d3;
    sch-drop-med-low_4 default-value d2;
    sch-drop-med-high_4 default-value d1;
    sch-drop-high_4 default-value d0;
    sch-drop-any_4 default-value d3;
    scheduler-map default-value all_smap;
  }
}

```



```

class-of-service {
  scheduler-maps {
    "$scheduler-map" {
      forwarding-class "$fc_1" scheduler "$sch_1";
      forwarding-class "$fc_2" scheduler "$sch_2";
      forwarding-class "$fc_3" scheduler "$sch_3";
      forwarding-class "$fc_4" scheduler "$sch_4";
    }
  }
  schedulers {
    "$sch_1" {
      transmit-rate "$sch-tx_1";
      buffer-size percent "$sch-bs_1";
      priority "$sch-pri_1";
      drop-profile-map loss-priority low protocol any drop-profile "$sch-drop-low_1";
      drop-profile-map loss-priority medium-low protocol any drop-profile "$sch-drop-med-low_1";
      drop-profile-map loss-priority medium-high protocol any drop-profile "$sch-drop-med-high_1";
      drop-profile-map loss-priority high protocol any drop-profile "$sch-drop-high_1";
    }
    "$sch_2" {
      transmit-rate percent "$sch-tx_2";
      buffer-size percent "$sch-bs_2";
      priority "$sch-pri_2";
      drop-profile-map loss-priority low protocol any drop-profile "$sch-drop-low_2";
      drop-profile-map loss-priority medium-low protocol any drop-profile "$sch-drop-med-low_2";
      drop-profile-map loss-priority medium-high protocol any drop-profile "$sch-drop-med-high_2";
      drop-profile-map loss-priority high protocol any drop-profile "$sch-drop-high_2";
    }
    "$sch_3" {
      transmit-rate "$sch-tx_3";
      buffer-size percent "$sch-bs_3";
      priority "$sch-pri_3";
      drop-profile-map loss-priority low protocol any drop-profile "$sch-drop-low_3";
      drop-profile-map loss-priority medium-low protocol any drop-profile "$sch-drop-med-low_3";
      drop-profile-map loss-priority medium-high protocol any drop-profile "$sch-drop-med-high_3";
      drop-profile-map loss-priority high protocol any drop-profile "$sch-drop-high_3";
    }
    "$sch_4" {
      transmit-rate percent "$sch-tx_4";
      buffer-size percent "$sch-bs_4";
      priority "$sch-pri_4";
      drop-profile-map loss-priority low protocol any drop-profile "$sch-drop-low_4";
      drop-profile-map loss-priority medium-low protocol any drop-profile "$sch-drop-med-low_4";
      drop-profile-map loss-priority medium-high protocol any drop-profile "$sch-drop-med-high_4";
    }
  }
}

```

```

        drop-profile-map loss-priority high protocol any drop-profile "$sch-drop-high_4";
    }
}
}
}
service_2 {
    variables {
        fc_1 default-value be;
        sch_1 default-value be_sch;
        sch-tx_1 default-value 10;
        sch-bs_1 default-value 10;
        sch-pri_1 default-value high;
        sch-drop-low_1 default-value d3;
        sch-drop-med-low_1 default-value d2;
        sch-drop-med-high_1 default-value d1;
        sch-drop-high_1 default-value d0;
        sch-drop-any_1 default-value d3;
        scheduler-map default-value all_smap;
    }
    class-of-service {
        scheduler-maps {
            "$scheduler-map" {
                forwarding-class "$fc_1" scheduler "$sch_1";
            }
        }
        schedulers {
            "$sch_1" {
                transmit-rate percent "$sch-tx_1";
                buffer-size percent "$sch-bs_1";
                priority "$sch-pri_1";
                drop-profile-map loss-priority low protocol any drop-profile "$sch-drop-low_1";
                drop-profile-map loss-priority medium-low protocol any drop-profile "$sch-drop-med-low_1";
                drop-profile-map loss-priority medium-high protocol any drop-profile "$sch-drop-med-high_1";
                drop-profile-map loss-priority high protocol any drop-profile "$sch-drop-high_1";
            }
        }
    }
}
}

```

4. The network administrator configures DHCP and RADIUS to grant access and services to the interfaces referenced by the **subscriber** dynamic profile.

```

[edit]
  forwarding-options {
    dhcp-relay {
      traceoptions {
        file size 1g;
        flag all;
      }
      dynamic-profile subscriber aggregate-clients replace;
      server-group {
        subscriber-server {
          203.0.113.2;
        }
      }
      active-server-group subscriber-server;
      group relay-0 {
        authentication {
          password $ABC123;
          username-include {
            user-prefix user0;
            mac-address;
          }
        }
        interface ge-1/1/0.100;
        interface ge-1/1/0.200;
      }
    }
  }
}
radius-server {
  198.51.100.11 secret "$ABC123$ABC123$ABC123"; ## SECRET-DATA
}
profile subscriber-profile {
  authentication-order radius;
  radius {
    authentication-server 198.51.100.11;
    accounting-server 198.51.100.11;
  }
  radius-server {
    198.51.100.11 secret "$ABC123$ABC123"; ## SECRET-DATA
  }
  accounting {
    order radius;
    statistics time;
  }
}

```

RELATED DOCUMENTATION

| [Configuring Per-Unit Scheduling in a Dynamic Profile](#) | 86

Configuring Dedicated Queue Scaling with Hierarchical CoS or Per-Unit Scheduling

IN THIS CHAPTER

- [Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview | 101](#)
- [Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces | 104](#)
- [Verifying the Number of Dedicated Queues Configured on MIC and MPC Interfaces | 106](#)

Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview

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. These queuing MPCs enable service providers to reduce the cost per subscriber by allowing many 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 MPCs

Beginning with Junos OS Release 15.1, MPC2E-3D-NG-Q, MPC3E-3D-NG-Q, MPC5EQ-40G10G, and MPC5EQ-100G10G MPCs support up to five levels of hierarchical queuing. Beginning with Junos OS Release 16.1R1, MPC7 line cards also support five levels of hierarchical queuing. [Table 13 on page 102](#) lists the number of dedicated queues and nodes supported per MPC.

Table 13: Dedicated Queues for MPCs

MPC	Dedicated Queues	Level 4 Nodes	Level 3 Nodes	Level 2 Nodes	Level 1 Nodes (Ports)
MPC2E-3D-NG-Q	512,000	64,000	16,000	4000	384
MPC3E-3D-NG-Q					
MPC5EQ-40G10G	1 million	128,000	32,000	4000	384
MPC5EQ-100G10G					
MPC7	256,000	32,000	8000	4000	126



CAUTION: The maximum scaling targets provided in [Table 13 on page 102](#) are based on system level design specifications. Actual realized subscriber or session scale is highly dependent upon the configuration and can be influenced by configuration variables including: the number of routes, the number of enabled services, the number of policy and firewall filters, policers, counters, statistics and access model type. Once you define a configuration, your Juniper account team can help characterize the expected system level scale or scale range for your live deployment.

MPCs vary in the number of Packet Forwarding Engines on board. MPC2E-3D-NG-Q and MPC3E-3D-NG-Q MPCs each have one Packet Forwarding Engine, allowing all 64,000 level 4 (subscriber) nodes to be allocated to a single MIC. MPC5EQ MPCs have two Packet Forwarding Engines, one for each possible MIC, each supporting 64,000 level 4 (subscriber) nodes.

NOTE: The nonqueuing MPCs MPC2E-3D-NG, MPC3E-3D-NG, MPC5E-40G10G, and MPC5E-100G10G provide up to eight queues per port in standard configuration. However, each of these MPCs can be configured to provide limited-scale hierarchical class of service (HCoS) and up to 32,000 queues.

Managing Remaining Queues

In Junos OS releases earlier than Release 15.1R4, SNMP traps generate system log messages to notify you:

- When the number of available dedicated queues on the MPC drops below 10 percent. For example:

```
Mar 15 14:55:22.977 host cosd[1963]: COSD_OUT_OF_DEDICATED_QUEUES: Queue usage
count for interface xe-3/0/0 is at 90 percent
```

- When the maximum number of dedicated queues on the MPCs is reached. For example,

```
Mar 15 18:01:59.344 host cosd[3848]: COSD_OUT_OF_DEDICATED_QUEUES: Queue usage
count for interface xe-3/0/0 is at 100 percent.
```

When the maximum number of dedicated queues is allocated, 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 is 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. These subscriber interfaces remain with this traffic-control profile, even if dedicated queues become available.

NOTE: Starting in Junos OS Release 15.1R4, the COSD_OUT_OF_DEDICATED_QUEUES functionality is not available for QoS-enabled dynamic subscribers. Starting in Junos OS Release 17.4R1, CoS resource monitoring enables you to set a per-FPC queue threshold of up to 90 percent of resources bound to a scheduling hierarchy; subscriber logins are not allowed when the threshold is reached. However, this threshold applies to all queues, not dedicated queues alone. See *Resource Monitoring for Subscriber Management and Services Overview* for more information.

Release History Table

Release	Description
16.1R1	Beginning with Junos OS Release 16.1R1, MPC7 line cards also support five levels of hierarchical queuing.
15.1R1	Beginning with Junos OS Release 15.1, MPC2E-3D-NG-Q, MPC3E-3D-NG-Q, MPC5EQ-40G10G, and MPC5EQ-100G10G MPCs support up to five levels of hierarchical queuing.

RELATED DOCUMENTATION

[Hierarchical Class of Service User Guide](#)

[Understanding Hierarchical Scheduling](#)

[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 | 104](#)

[Understanding Hierarchical Scheduling for MIC and MPC Interfaces](#)

Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces

IN THIS SECTION

- [Configuring the Maximum Number of Queues for MIC and MPC Interfaces | 104](#)
- [Configuring Remaining Common Queues on MIC and MPC Interfaces | 105](#)

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

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# set max-queues-per-interface (8 | 4)
```

SEE ALSO

[Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview](#) | 101

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
```

SEE ALSO

[Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview | 101](#)

RELATED DOCUMENTATION

[Verifying the Number of Dedicated Queues Configured on MIC and MPC Interfaces | 106](#)

[Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview | 101](#)

[Configuring Static Hierarchical Scheduling in a Dynamic Profile](#)

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 | 104](#)

Shaping Downstream Traffic Based on Frames or Cells

IN THIS CHAPTER

- [Bandwidth Management for Downstream Traffic in Edge Networks Overview | 109](#)
- [Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 111](#)
- [Example: Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 113](#)
- [Configuring Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 117](#)
- [Example: Configuring Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 118](#)
- [Setting Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags | 121](#)
- [Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces | 123](#)
- [Reporting the Effective Shaping Rate for Subscribers | 124](#)
- [Verifying the Effective Shaping Rate Reporting Configuration | 124](#)

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 downstream network is not necessarily the directly attached network. In typical broadband network gateway (BNG) configurations, the directly attached network is an Ethernet access network, which provides access to either another frame-based network, or a cell-based network.

The *overhead accounting* feature enables you to shape traffic based on whether the downstream network is a frame-based network, like Ethernet, or a cell-based network, like ATM. It assigns a byte adjustment value to account for different encapsulations.

This feature is available on MIC and MPC interfaces.

Effective Shaping Rate

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.
- Number of bytes in each frame that are accounted for by the shaper.

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.

Shaping Modes

There are two modes used for adjusting downstream traffic:

- *Frame shaping mode* is useful for adjusting downstream traffic with different encapsulations. 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 router adjusts all of the rates by 48/53 to account for 5-byte 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.

Byte Adjustments

When the downstream traffic has different byte sizes per encapsulation, it is useful to configure a *byte adjustment* value to adjust the number of bytes per packet to be included in or excluded from the shaping mechanism. This value represents the number of bytes that are encapsulated and decapsulated by the downstream equipment. 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.

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, 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.

RELATED DOCUMENTATION

To configure overhead accounting for static Ethernet interfaces, see [Configuring Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 117](#)

To configure overhead accounting for dynamic subscriber access, see [Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 111](#)

[Setting Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags | 121](#)

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 | 109](#)

[Example: Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 113](#)

[Verifying the Scheduling and Shaping Configuration for Subscriber Access | 66](#)

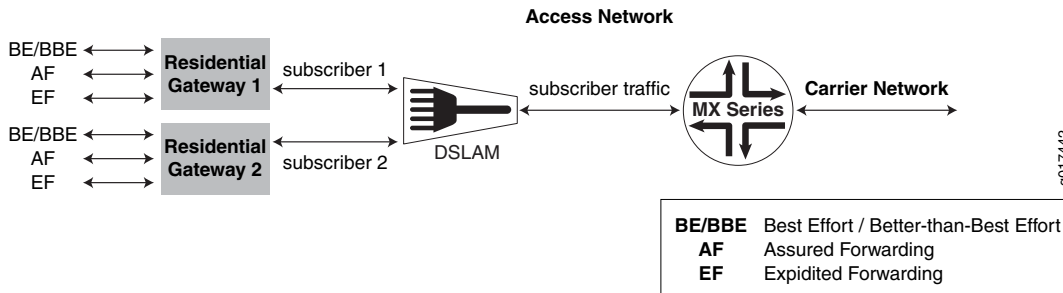
Example: Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates

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 2 on page 113 shows the sample network that the examples reference.

Figure 2: Sample Network Topology for Downstream Traffic



Managing Traffic with Different Encapsulations

In this example, the MX Series router shown in Figure 2 on page 113 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.

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 14 on page 114](#) lists the initial values defined by the RADIUS administrator for the shaping rates.

Table 14: Initial Shaping Values at Subscriber Login For Traffic With Different Encapsulations

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

Table 14: Initial Shaping Values at Subscriber Login For Traffic With Different Encapsulations (*continued*)

Predefined Variable	RADIUS Tag	Value
\$junos-cos-byte-adjust	T08	-4

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 2 on page 113](#) 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 15 on page 116 lists the initial values defined by the RADIUS administrator for the shaping rates.

Table 15: Initial Shaping Values at Subscriber Login For Downstream Cell-Based Traffic

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](#) | 111

Configuring Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates

The overhead accounting feature enables you to account for downstream traffic that has different encapsulations or downstream traffic from cell-based equipment, such as ATM switches.

You can configure the overhead accounting feature to shape downstream traffic based on frames or cell shaping mode.

You can also account for the different byte sizes per encapsulation by configuring a byte adjustment value for the shaping mode.

To configure the shaping mode and byte adjustment value for static CoS configurations:

1. Specify the shaping mode.

Frame shaping mode is enabled by default.

```
[edit class-of-service traffic-control-profiles profile-name]
user@host# set overhead-accounting (frame-mode | cell-mode)
```

2. (Optional) Specify a byte adjustment value.

```
[edit class-of-service traffic-control-profiles profile-name]
user@host# set overhead-accounting bytes byte-value
```

BEST PRACTICE: We recommend that you specify a byte adjustment value that represents the difference between the customer premise equipment (CPE) protocol overhead and the 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

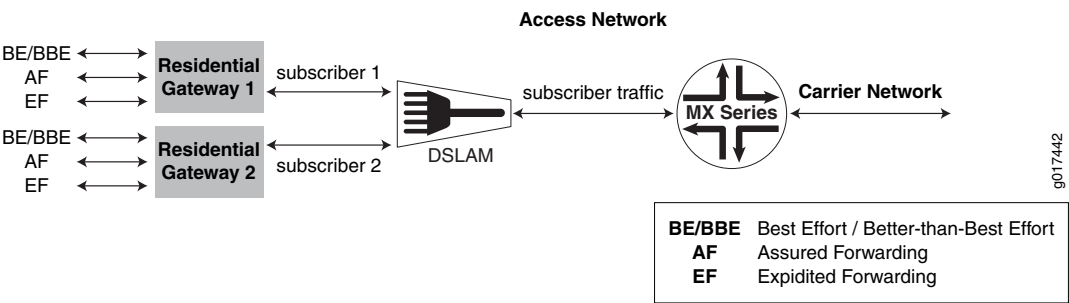
| [Bandwidth Management for Downstream Traffic in Edge Networks Overview](#) | 109

Example: Configuring Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates

This topic describes two scenarios for which you can configure static shaping parameters to account for packet overhead in a downstream network.

[Figure 2 on page 113](#) shows the sample network that the examples reference.

Figure 3: Sample Network Topology for Downstream Traffic



Managing Traffic with Different Encapsulations

In this example, the MX Series router shown in [Figure 2 on page 113](#) 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.

To account for the different frame sizes, the network administrator configures the frame shaping mode with -4 byte adjustment:

1. The network administrator configure the traffic shaping parameters and attaches 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]
class-of-service {
  traffic-control-profiles {
    tcp-example-overhead-accounting-frame-mode {
      shaping-rate 10m;
      shaping-rate-priority-high 4m;
      guaranteed-rate 2m;
      excess-rate percent 50;
      overhead-accounting frame-mode bytes -4;
    }
  }
  interfaces {
    ge-1/0/0 {
      output-traffic-control-profile tcp-example-overhead-accounting-frame-mode;
    }
  }
}
```

2. The network administrator verifies the adjusted rates.

user@host#**show class-of-service traffic-control-profile**

```
Traffic control profile: tcp-example-overhead-accounting-frame-mode, Index:
61785
Shaping rate: 10000000
Shaping rate priority high: 4000000
Excess rate 50
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 2 on page 113](#) 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.

To account for the different frame sizes, the network administrator configures the cell shaping mode with -4 byte adjustment:

1. Configure the traffic shaping parameters 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]
class-of-service {
  traffic-control-profiles {
    tcp-example-overhead-accounting-cell-mode {
      shaping-rate 10m;
      shaping-rate-priority-high 4m;
      guaranteed-rate 2m;
      excess-rate percent 50;
      overhead-accounting cell-mode;
    }
  }
  interfaces {
    ge-1/0/0 {
      output-traffic-control-profile tcp-example-overhead-accounting-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
Shaping rate priority high: 4000000
Excess rate 50
Guaranteed rate: 2000000
```

```
Overhead accounting mode: 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 Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates](#) | 117

Setting Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags

You can use access line parameters received 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 **dynamic-class-of-service-options** statement at the **[edit dynamic-profiles *profile-name* class-of-service]** hierarchy level. Specify the appropriate attribute as a value for the **vendor-specific-tags** option.

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-minimum-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](#) | 109

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 Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags | 121](#)

[Bandwidth Management for Downstream Traffic in Edge Networks Overview | 109](#)

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 | 124](#)

Hierarchical CoS Shaping-Rate Adjustments Overview

[Bandwidth Management for Downstream Traffic in Edge Networks Overview | 109](#)

Juniper Networks VSAs Supported by the AAA Service Framework

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

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;
...
```

- 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](#) | 124

Applying CoS to Households or Individual Subscribers Using ACI-Based Dynamic VLANs

IN THIS CHAPTER

- [Applying CoS Attributes to VLANs Using Agent-Circuit-Identifiers | 127](#)
- [Agent Circuit Identifier-Based Dynamic VLANs Bandwidth Management Overview | 130](#)
- [Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ACI Interface Sets | 134](#)
- [Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs | 135](#)

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 ACI VLAN interface set using the dynamic profile for the ACI 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 interface set using a unique-ID based dynamic scheduler map:

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
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:

```
dynamic-profiles {
  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 {
        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](#) | 40

[Changing CoS Services Overview](#) | 160

Agent Circuit Identifier-Based Dynamic VLANs Bandwidth Management Overview

IN THIS SECTION

- [CoS Shaping Rate Adjustment](#) | 131
- [CoS Overhead Accounting Adjustment](#) | 131
- [Dynamic Profiles and Adjustment of CoS Shaping Rate and Overhead Accounting](#) | 132
- [Guidelines for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting](#) | 133

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

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 kilobits 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.

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.

[Table 16 on page 133](#) 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 16: CoS Adjustment in Dynamic Profiles for ACI Interface Sets and ACI-Based Subscriber Interfaces

VSA's 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 134](#), 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 Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags | 121](#)

[Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs | 135](#)

[Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ACI Interface Sets | 134](#)

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 | 130](#)

[Setting Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags | 121](#)

[Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Agent Circuit Identifier-Based Dynamic VLANs | 135](#)

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 ACI Interface Sets*.

- 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*.

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 | 130](#)

[Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ACI Interface Sets | 134](#)

Configuring Dynamic VLANs Based on Agent Circuit Identifier Information

Applying CoS to Households or Individual Subscribers Using Access Line Identifier Dynamic VLANs

IN THIS CHAPTER

- Applying CoS Attributes to VLANs Using Access-Line Identifiers | 137
- Bandwidth Management Overview for Dynamic VLANs Based on Access-Line Identifiers | 140
- Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ALI Interface Sets | 144
- Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Dynamic VLANs Based on Access-Line Identifiers | 145

Applying CoS Attributes to VLANs Using Access-Line Identifiers

To apply CoS attributes, such as shaping, at the household level, you must set and define the CoS policy for the access-line-identifier (ALI) VLAN interface set using the dynamic profile for the ALI 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 ALI interface set using a unique-ID based dynamic scheduler map:

Configure a CoS dynamic profile with a simple traffic-control profile that is applied to the dynamic interface set that represents the ALI 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
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 ALI set using a unique ID-based dynamic scheduler map:

```
dynamic-profiles {
  ali-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 {
        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](#) | 40

[Changing CoS Services Overview](#) | 160

Bandwidth Management Overview for Dynamic VLANs Based on Access-Line Identifiers

IN THIS SECTION

- [CoS Shaping Rate Adjustment](#) | 141
- [CoS Overhead Accounting Adjustment](#) | 141
- [Dynamic Profiles and Adjustment of CoS Shaping Rate and Overhead Accounting](#) | 142
- [Guidelines for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting](#) | 143

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 access-line identifiers (ALI), you can shape the bandwidth either at a per-household level for a dynamic ALI interface set, or at a per-subscriber level for a dynamic VLAN subscriber interface associated with an ALI interface set.

To help you manage bandwidth efficiently and economically for ALI-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 ALI interface sets and their associated ALI-based dynamic VLAN subscriber interfaces.

This overview covers the following topics:

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 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 ALI 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.

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 DSL 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 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 ALI 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 either 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 ALI interface set is created and the router receives the PADI and PADR packets from the first subscriber interface belonging to the ALI 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 ALI 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 ALI-based dynamic PPPoE (**pp0**) subscriber interface associated with the ALI interface set.

[Table 17 on page 142](#) summarizes how the dynamic profile in which you configure CoS adjustment for ALI-based dynamic VLANs using one or both VSAs affects the router behavior.

Table 17: CoS Adjustment in Dynamic Profiles for ALI Interface Sets and ALI-Based Subscriber Interfaces

VSA Specified in ALI Interface Set Dynamic Profile	VSA Specified in PPPoE Subscriber Interface Dynamic Profile	Result
Yes	No	Router adjusts specified CoS attributes only for dynamic ALI interface set

Table 17: CoS Adjustment in Dynamic Profiles for ALI Interface Sets and ALI-Based Subscriber Interfaces *(continued)*

VSA's Specified in ALI Interface Set Dynamic Profile	VSA's Specified in PPPoE Subscriber Interface Dynamic Profile	Result
No	Yes	Router adjusts specified CoS attributes only for ALI-based dynamic PPPoE subscriber interface
Yes	Yes	Router adjusts specified CoS attributes for both dynamic ALI interface set and ALI-based dynamic PPPoE subscriber interface
No	No	Router does not adjust CoS attributes for either the dynamic ALI interface set or the ALI-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 ALI interface sets.

With the exception of the constraints described in [“Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ALI Interface Sets”](#) on page 144, most of the guidelines and restrictions that apply to this feature for use with dynamic subscriber interfaces that are not based on ALIs also apply to its use for dynamic ALI interface sets and their associated ALI-based dynamic VLAN subscriber interfaces.

RELATED DOCUMENTATION

[Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ALI Interface Sets](#) | 144

[Setting Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags](#) | 121

[Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Dynamic VLANs Based on Access-Line Identifiers](#) | 145

Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic ALI 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 interface sets based on the access-line identifier (ALI) and their associated ALI-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

[Bandwidth Management Overview for Dynamic VLANs Based on Access-Line Identifiers](#) | 140

[Setting Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags](#) | 121

[Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Dynamic VLANs Based on Access-Line Identifiers](#) | 145

Adjusting the CoS Shaping Rate and Overhead Accounting Parameters for Dynamic VLANs Based on Access-Line Identifiers

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 access-line-identifier (ALI) interface sets and their associated ALI-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 ALI interface set.

See *Defining Access-Line-Identifier Interface Sets*.

- To configure adjustment of the CoS shaping rate and overhead accounting attributes on a per-subscriber basis, create a dynamic profile that defines the ALI-based dynamic PPPoE (pp0) subscriber interface associated with the ALI interface set.

See *Configuring Dynamic VLAN Subscriber Interfaces Based on Access-Line Identifiers*.

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 ALI interface sets and associated ALI-based dynamic VLAN subscriber interfaces, do either or both of the following:

- In a dynamic profile for an ALI interface set or a dynamic profile for an ALI-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 ALI interface set or a dynamic profile for an ALI-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

Configuring Dynamic VLANs Based on Access-Line Identifiers

[Bandwidth Management Overview for Dynamic VLANs Based on Access-Line Identifiers | 140](#)

[Restrictions for Configuring Adjustment of CoS Shaping Rate and Overhead Accounting for Dynamic
All Interface Sets | 144](#)

Managing Excess Bandwidth Distribution and Traffic Bursts

IN THIS CHAPTER

- [Excess Bandwidth Distribution on MIC and MPC Interfaces Overview | 147](#)
- [Traffic Burst Management on MIC and MPC Interfaces Overview | 148](#)
- [Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces | 151](#)

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 MPC Interfaces | 151](#)

Per-Priority Shaping on MIC and MPC Interfaces Overview

[Traffic Burst Management on MIC and MPC Interfaces Overview | 148](#)


Traffic Burst Management on MIC and MPC Interfaces Overview

IN THIS SECTION

- Guidelines for Configuring the Burst Size | 148
- How the System Calculates the Burst Size | 150

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 4: Sample Burst Shaping Rates



In [Figure 4 on page 148](#), 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

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.

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 150](#).

- 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. Starting in Junos OS Release 15.1, the CLI no longer generates a commit error when the guaranteed-rate burst size is statically configured to be more than the shaping-rate burst size. This behavior changed with the advent of enhanced subscriber management. The system logs an error when the guaranteed-burst rate is higher, whether it is configured statically, dynamically with predefined variables, or by means of a change of authorization request.
- 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.
- Burst-size is not supported with **per-priority-shaping**.

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.

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, the CLI no longer generates a commit error when the guaranteed-rate burst size is statically configured to be more than the shaping-rate burst size.

RELATED DOCUMENTATION

Per-Priority Shaping on MIC and MPC Interfaces Overview
Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC 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 )
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](#) | 40

Applying CoS Using Parameters Received from RADIUS

IN THIS CHAPTER

- [Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS | 155](#)
- [Changing CoS Services Overview | 160](#)
- [CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview | 164](#)
- [Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions | 166](#)
- [Configuring Initial CoS Parameters Dynamically Obtained from RADIUS | 167](#)
- [Configuring Static Default Values for Traffic Scheduling and Shaping | 168](#)
- [Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions | 170](#)
- [CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets | 172](#)
- [Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS | 178](#)

Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS

IN THIS SECTION

- [Dynamic Configuration of Initial CoS in Access Profiles | 156](#)
- [Predefined Variables for Dynamic Configuration of Initial Traffic Shaping | 156](#)
- [Predefined Variables for Dynamic Configuration of Initial Scheduling and Queuing | 157](#)

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 5G Universal Routing Platforms.

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

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 CoS 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 18 on page 157](#) in an access dynamic profile associated with the subscriber interface.

Table 18: CoS Predefined Variables for Scheduler Map and Traffic Shaping

Variable	Description
\$junos-cos-scheduler-map	<p>Scheduler-map name to be dynamically configured in a traffic-control profile in the access dynamic profile when a subscriber logs in.</p> <p>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).</p>
\$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.

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 19 on page 158](#) in an access dynamic profile associated with the subscriber interface.

Table 19: 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	<p>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.</p> <p>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.</p>

Table 19: CoS Predefined Variables for Scheduling and Queuing (*continued*)

Variable	Description
\$junos-cos-scheduler-dropfile-medium-low	<p>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.</p> <p>NOTE: The drop profile must be configured statically (at the [edit class-of-service schedulers scheduler-name drop-profiles] hierarchy level).</p>
\$junos-cos-scheduler-dropfile-medium-high	<p>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.</p> <p>NOTE: The drop profile must be configured statically (at the [edit class-of-service schedulers scheduler-name drop-profiles] hierarchy level).</p>
\$junos-cos-scheduler-dropfile-high	<p>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.</p> <p>NOTE: The drop profile must be configured statically (at the [edit class-of-service schedulers scheduler-name drop-profiles] hierarchy level).</p>
\$junos-cos-scheduler-dropfile-any	<p>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.</p> <p>NOTE: The drop profile must be configured statically (at the [edit class-of-service schedulers scheduler-name drop-profiles] hierarchy level).</p>

RELATED DOCUMENTATION

Subscriber Activation and Service Management in an Access Network

Dynamic Profiles Overview

Dynamic Variables Overview

Junos OS Predefined Variables

[Configuring Initial CoS Parameters Dynamically Obtained from RADIUS | 167](#)

[Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS | 178](#)

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 a *dynamic client 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 client profile with Junos OS predefined variables for RADIUS-provided CoS parameters, you also enable the service to be activated for those subscribers at login.

NOTE: The dynamic client profile is also referred to as a dynamic client access profile, or sometimes just access profile for brevity. Do not confuse this profile, configured at the **[edit dynamic-profiles profile-name]** hierarchy level, with the access profile configured at the **[edit access profile profile-name]** hierarchy level. These static access profiles are used to configure authentication, accounting, and authorization parameters for subscriber access, some session attributes, and client-specific properties for L2TP and PPP sessions. Access profiles are applied at various configuration levels with the **access-profile** statement.

To enable subscribers to activate a service or upgrade to different services through RADIUS change-of-authorization (CoA) messages after login, configure a *dynamic 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 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 20 on page 162](#) 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 20: 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	<ul style="list-style-type: none"> 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 	<ul style="list-style-type: none"> Configure RADIUS values or default values for all parameters in access profile Configure scheduler map and schedulers in access profile 	<ul style="list-style-type: none"> Configure RADIUS values or default values for all parameters in access profile Configure scheduler map and schedulers in access profile 	<ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> Delay buffer rate Guaranteed rate Scheduler map Shaping rate 	Replaces the following parameters: <ul style="list-style-type: none"> Delay buffer rate Guaranteed rate Shaping rate Scheduler map 	Combines the values of the following parameters to their maximum scalar value: <ul style="list-style-type: none"> Delay buffer rate Guaranteed rate Shaping rate Appends the scheduler map parameter	Replaces the following parameters: <ul style="list-style-type: none"> Delay buffer rate Guaranteed rate Shaping rate Scheduler map

Table 20: 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	<p>Does not merge queues</p> <p>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.</p>	<p>Merge queues if the queue specified in the service profile is not already in use for the subscriber</p> <p>NOTE: Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.</p>	<p>Merge queues if the queue specified in the service profile is not already in use for the subscriber</p> <p>NOTE: Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.</p>	<p>Merge queues if the queue specified in the service profile is not already in use for the subscriber</p> <p>NOTE: Do not instantiate a CoA request using a service dynamic profile that is already in use on the same logical interface.</p>

RELATED DOCUMENTATION

[Configuring Static Hierarchical Scheduling in a Dynamic Profile](#)

[Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview](#)

[RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework](#)

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview

IN THIS SECTION

- [Supported Network Configurations | 164](#)
- [Traffic-Control Profiles in Subscriber Interface Dynamic Profiles | 164](#)
- [CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets and Member Subscriber Sessions | 165](#)

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

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

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 172](#).

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 | 166](#)

[Applying CoS Traffic-Shaping Attributes to Dynamic Interface Sets and Member Subscriber Sessions | 170](#)

[CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets | 172](#)

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 172](#).

- 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 | 170](#)

[CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets | 172](#)

[CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions](#)

[Overview | 164](#)

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 on certain subscriber interfaces, configure Juniper Networks VSA 28–146.

See *Configuring Router or Switch Interaction with RADIUS Servers* and *RADIUS Servers and Parameters for Subscriber Access*.

2. Configure a subscriber interface that supports hierarchical CoS.
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 217](#).
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 49](#).
5. Configure forwarding classes and scheduler maps statically.

See *Configuring a Custom Forwarding Class for Each Queue* 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 54](#).

RELATED DOCUMENTATION

[Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS | 155](#)

[Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS | 178](#)

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[Subscriber Activation and Service Management in an Access Network](#)

[Juniper Networks VSAs Supported by the AAA Service Framework](#)

[Dynamic Profiles Overview](#)

[Dynamic Variables Overview](#)

[Junos OS Predefined Variables](#)

Configuring Static Default Values for Traffic Scheduling and Shaping

To provide subscribers with default values for CoS parameters, configure user-defined variables for CoS parameters and assign static default values to the variables. If you have configured values to be supplied by a RADIUS CoA, subscribers receive the default value when deactivating a service.

To configure user-defined variables with default values 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. Configure a default value for the shaping rate.

```
[edit dynamic-profiles residential-silver variables]
user@host# set srate default-value 5m
```

3. Configure a default value for the guaranteed rate.

```
[edit dynamic-profiles residential-silver variables]
user@host# set grate default-value 5m
```

4. Configure a default value for the delay buffer rate.

```
[edit dynamic-profiles residential-silver variables]
user@host# set dbrate default-value 10m
```

5. 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. Access the traffic-control profile 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"
```

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access](#) | 40

[Changing CoS Services Overview](#) | 160

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-identifier (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](#)
- [Configuring a Dynamic Profile Used to Create Single-Tag VLANs](#)
- [Configuring a Dynamic Profile Used to Create Stacked VLANs](#)
- [Configuring Dynamic PPPoE Subscriber Interfaces](#)
- [Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles](#)

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 172](#) 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 {
    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 | 172](#)

[CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview | 164](#)

[Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions | 166](#)

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 21 on page 173](#).

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 21 on page 173](#). 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 21 on page 173](#) 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 21 on page 173](#) use RADIUS vendor ID 4874 and RADIUS attribute value 108.

Table 21: 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.

Table 21: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets (continued)

Predefined Variable	RADIUS Tag Value	Description
\$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.
\$junos-cos-excess-rate	105	<p>Excess rate configured in a traffic-control profile in a dynamic profile; scheduler weighting when operating in the excess region between the guaranteed rate and the shaping rate.</p> <p>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.</p>
\$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	<p>Byte adjustment value for the cell or frame shaping mode configured in a traffic-control profile in a dynamic profile.</p> <p>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.</p>

Table 21: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets (*continued*)

Predefined Variable	RADIUS Tag Value	Description
\$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.
\$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.

Table 21: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets (*continued*)

Predefined Variable	RADIUS Tag Value	Description
\$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.
\$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.

Table 21: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets (*continued*)

Predefined Variable	RADIUS Tag Value	Description
\$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.
\$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.

Table 21: Junos OS CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets (continued)

Predefined Variable	RADIUS Tag Value	Description
\$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 | 170](#)

[CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview | 164](#)

[Guidelines for Configuring CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions | 166](#)

Junos OS Predefined Variables

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;
                    }
                }
            }
        }
    }
}

```

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

Dynamic Profiles Overview

Dynamic Variables Overview

Junos OS Predefined Variables

[Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS | 155](#)

[Configuring Initial CoS Parameters Dynamically Obtained from RADIUS | 167](#)

Modifying a Subscriber’s Shaping Characteristics After a Subscriber is Instantiated

IN THIS CHAPTER

- CoS Adjustment Control Profiles Overview | 183
- Configuring CoS Adjustment Control Profiles | 186
- Verifying the CoS Adjustment Control Profile Configuration | 188

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, DHCP tags, and RADIUS Change of Authorization (CoA) messages after a subscriber is instantiated. Adjustment control profiles are router-wide and apply to both static and dynamic interfaces.

[Table 22 on page 183](#) describes the applications and their associated default algorithms that can be configured to perform rate adjustments after the subscriber is instantiated.

Table 22: 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 22: 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.
DHCP	2	Adjust-less	<p>The DHCP application can include DSL Forum VSA attributes in its discovery messages, DHCPDISCOVER for DHCPv4 and SOLICIT for DHCPv6.</p> <p>The attributes can modify the Junos OS CLI-configured shaping-rate value, as well as the RADIUS-supplied shaping-rate value. By default, these values can be modified by subsequent RADIUS CoA messages and DHCP actions.</p> <p>The DSL Forum VSAs are conveyed in DHCP option 82, suboption 9 (Vendor-Specific Information suboption) for DHCPv4 and in Option 17 (Vendor-Specific Information option) for DHCPv6.</p>
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-supplied 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.

The lower the priority value, the higher the priority. For example priority 1 is higher than priority 2. By default, the application shaping rates compare as follows:

- ANCP has priority over all the other applications.
- RADIUS CoA has priority over DHCP tags or PPPoE IA tags.
- The DHCP tags or PPPoE IA tags have priority over the shaping rate configured in the traffic control profile.

Applications and Associated Algorithms in Adjustment Control Profiles

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:

NOTE: All values can apply to shaping rates. Only adjust-never and adjust-always can apply to overhead-accounting attributes.

- 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.
- 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.

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.

CoS Shaping Rate Fallback Behavior

When a CoS service profile is deactivated or removed, the CoS shaping rate falls back to the next highest available adjustment source as follows:

1. Fall back to the ANCP shaping rate if it is present and it has a higher priority than RADIUS CoA, the DHCP tags, or the PPPoE IA tags.
2. Fall back to the RADIUS CoA shaping rate if it is present and it has a higher priority than the DHCP tags or the PPPoE IA tags.
3. Fall back to the DHCP tags or the PPPoE IA tags shaping rate, if present.
4. Fall back to the shaping rate configured in the associated traffic control profile.

When a shaping rate is adjusted by ANCP, if that adjustment is removed, the rate reverts to the PPPoE IA tag rate if it is present. If the tag rate is not present then the shaping rate reverts to the configured rate in the traffic control profile.

When an ANCP adjustment for overhead-accounting mode is removed, the value reverts to the PPPoE IA tag value if it is present. If the tag value is not present, then the mode reverts to the configured value in the traffic control profile.

When an ANCP adjustment for overhead-accounting bytes is removed, the value reverts to the configured value in the traffic control profile; PPPoE IA tags cannot provide this value.

RELATED DOCUMENTATION

[Configuring CoS Adjustment Control Profiles | 186](#)

[Verifying the CoS Adjustment Control Profile Configuration | 188](#)

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) Configure the adjustment controls for the DHCP application.

```
[edit class-of-service adjustment-control-profiles profile-name ]
user@host# set application dhcp-tags priority priority algorithm algorithm
```

6. (Optional) Verify your configuration.

```
user@host> show class-of-service adjustment-control-profiles
```

```
name: ANCP, priority: 1, algorithm: less;
name: RADIUS CoA, priority: 1, algorithm: always;
name: PPPoE IA tags, priority: 2, algorithm: less;
name: DHCP tags, priority: 2, algorithm: less
```

RELATED DOCUMENTATION

[CoS Adjustment Control Profiles Overview | 183](#)

[Verifying the CoS Adjustment Control Profile Configuration | 188](#)

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
name: DHCP tags, priority: 2, algorithm: less
```

RELATED DOCUMENTATION

[CoS Adjustment Control Profiles Overview | 183](#)

[Configuring CoS Adjustment Control Profiles | 186](#)

[adjustment-control-profiles | 712](#)

[application \(Adjustment Control Profiles\) | 729](#)

Applying CoS to Groups of Subscriber Interfaces

IN THIS CHAPTER

- [CoS for Interface Sets of Subscribers Overview | 189](#)
- [Configuring an Interface Set of Subscribers in a Dynamic Profile | 192](#)
- [Example: Configuring a Dynamic Interface Set of VLAN Subscribers | 193](#)
- [Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile | 211](#)

CoS for Interface Sets of Subscribers Overview

IN THIS SECTION

- [Guidelines for Configuring Dynamic Interface Sets in a Subscriber Access Network | 189](#)

Interface sets enable service providers to group logical interfaces or other interface sets 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

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.

- 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-aggregation-interface-set-name** is the L2 interface-set representing a logical intermediate node (DPU-C or PON tree) in the access network.
- The **\$junos-phy-ifd-underlying-intf-set-name** represents a default, topology-based interface-set (based on the physical interface name with a post-pend of “-underlying”) to conserve L2 CoS nodes.
- 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-phy-ifd-interface-set-name** predefined variable locally generates an interface set name associated with the underlying physical interface in a dynamic profile. This predefined variable enables you to group all the subscribers on a specific physical interface so that you can apply services to the entire group of subscribers.

Another use case for this predefined variable is to conserve CoS resources in a mixed business and residential topology by collecting the residential subscribers into an interface set associated with the physical interface, so that a level 2 node is used for the interface set rather than for each residential interface. Otherwise, because the business and residential subscribers share the same interface and business subscribers require three levels of CoS, then three levels are configured for each residential subscriber. That results in an unnecessary level 2 node being consumed for each residential connection, wasting CoS resources.

- 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.

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](#) | 192

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.

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
```

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](#) | 189

Guidelines for Configuring Dynamic CoS for Subscriber Access	40
CoS for Interface Sets of Subscribers Overview	189
Example: Configuring a Dynamic Interface Set of VLAN Subscribers	193
CoS for Aggregated Ethernet Subscriber Interfaces Overview	44

Example: Configuring a Dynamic Interface Set of VLAN Subscribers

IN THIS SECTION

- Requirements | 193
- Overview | 193
- Configuring the Dynamic VLANs | 193
- Configuring Dynamic Traffic Scheduling and Shaping | 196
- Configuring the Interface Set in the Dynamic Profile | 200
- Configuring DHCP Access | 202
- Configuring RADIUS Authentication | 204
- Verification | 210

Requirements

This example uses the following software and hardware components:

- 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]
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 203.0.113.32
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 203.0.113.32/32
```

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 203.0.113.32
```

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
```

- 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:

- Create the loopback interface.

```
[edit]
user@host# edit interfaces lo0
```

- Configure the unit and the family.

```
[edit interfaces lo0]
user@host# set unit 0 family inet address 203.0.113.32/32
```

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
```



```

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
```

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 203.0.113.32
```

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
```

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 $ABC123
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 203.0.113.0/16
set range limited low 203.0.113.10
set range limited high 203.0.113.250
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 $ABC123
```

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 203.0.113.0/16
```

3. Configure the upper and lower bounds of the address range.

```
[edit access address-assignment pool v4]
user@host#set range limited low 203.0.113.10
user@host#set range limited high 203.0.113.250
```

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 192.51.100.108
set secret $ABC123ABC123ABC123
set timeout 5
set retry 5
up 2
edit profile acc-prof
set authentication-order radius
set radius authentication-server 192.51.100.108
```

Configuring RADIUS Access

Step-by-Step Procedure

To configure RADIUS access:

1. Configure the RADIUS server.

```
[edit access]
user@host#edit radius-server 192.51.100.108
```


2. Configure the required secret (password) that the local router or switch passes to the RADIUS client.

```
[edit access radius-server 192.51.100.108]
user@host# set secret $ABC123ABC123ABC123
```

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 192.51.100.108]
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 192.51.100.108]
user@host# set retry 5
```

5. Configure the access profile.

```
[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 192.51.100.108
```

Results

```
dynamic-profiles {
  vlan-prof {
    interfaces {
```

```

    "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
            vlan-id "$junos-vlan-id";
            demux-source inet;
            family inet {
                unnumbered-address lo0.0 preferred-source-address 203.0.113.32;
            }
        }
    }
}
}
}
}
multiplay {
    class-of-service {
        traffic-control-profiles {
            multiplay {
                scheduler-map all_smap;
                shaping-rate 100m;
                guaranteed-rate 20m;
            }
        }
        interfaces {
            interface-set "$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 {
    192.51.100.108 {
      secret "$ABC123ABC123ABC123"; ## SECRET-DATA
      timeout 5;
      retry 5;
    }
  }
  profile acc-prof {
    authentication-order radius;
    radius {
      authentication-server 192.51.100.108;
    }
  }
  address-assignment {
    pool v4 {
      family inet {
        network 203.0.113.0/16;
        range limited {
          low 203.0.113.10;
          high 203.0.113.250;
        }
        dhcp-attributes {
          maximum-lease-time 84600;
        }
      }
    }
  }
}
class-of-service {
  interfaces {
    interface-set dynamic-set {
      output-traffic-control-profile multiplay;
    }
  }
}
interfaces {
  interface-set "$junos-interface-set-name" {
    interface "$junos-interface-ifd-name" {
      unit "$junos-underlying-interface-unit";
    }
  }
  "$junos-interface-ifd-name" {
    unit "$junos-underlying-interface-unit" {

```



```

    }
  }
}

```

Verification

IN THIS SECTION

- [Verifying the Interfaces that are Included in the Interface Set | 210](#)
- [Verifying the Traffic Scheduling and Shaping Parameters for the Interface Set | 210](#)

To confirm that the configuration is correct, perform these tasks:

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 Hierarchical CoS for Subscriber Interfaces](#)

[Configuring an Interface Set of Subscribers in a Dynamic Profile | 192](#)

Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile

IN THIS SECTION

- Requirements | 211
- Overview | 211
- Configuration | 212
- Verification | 215

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

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 *DHCP Subscriber Interface Overview* and in the following:

- For dynamic VLAN interfaces, see *Configuring a Static or Dynamic VLAN Subscriber Interface over Aggregated Ethernet*.
- For dynamic IP demux interfaces, see *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles* and *Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet*.
- For dynamic VLAN demux interfaces, see *Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles*.

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.
- **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]
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](#)

[Configuring a Basic Dynamic Profile](#)

[Configuring Hierarchical Schedulers for CoS](#)

[Configuring Remaining Common Queues on MIC and MPC Interfaces](#) | 105

Applying CoS to Subscriber Interfaces

IN THIS CHAPTER

- Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile | 217
- Applying Minimal Shaping and Scheduling to Remaining Subscriber Traffic | 218
- Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile | 219
- Applying a Classifier to a Subscriber Interface in a Dynamic Profile | 221

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
```

RELATED DOCUMENTATION

For hardware requirements and configuration guidelines, see [Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

Configuring Static Hierarchical Scheduling in a Dynamic Profile

Example: Maintaining a Constant Traffic Flow by Configuring a Static VLAN Interface with a Dynamic Profile for Subscriber Access

Example: Configuring Dynamic Hierarchical Scheduling for Subscribers

[Verifying the Scheduling and Shaping Configuration for Subscriber Access | 66](#)

[CoS for Subscriber Access Overview | 39](#)

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
```

RELATED DOCUMENTATION

[Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile](#) | 217

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.

```
[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)
```

RELATED DOCUMENTATION

For hardware requirements and configuration guidelines, see [Guidelines for Configuring Dynamic CoS for Subscriber Access](#) | 40

Example: Configuring Dynamic Hierarchical Scheduling for Subscribers

[Verifying the Scheduling and Shaping Configuration for Subscriber Access](#) | 66

[Applying a Classifier to a Subscriber Interface in a Dynamic Profile](#) | 221

Applying IEEE 802.1p Rewrite Rules to Dual VLAN Tags

Rewriting Packet Headers to Ensure Forwarding Behavior

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 *Configuring Behavior Aggregate 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

For hardware requirements and configuration guidelines, see [Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

Example: Configuring Dynamic Hierarchical Scheduling for Subscribers

[Verifying the Scheduling and Shaping Configuration for Subscriber Access | 66](#)

[Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile | 219](#)

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

3

PART

Configuring Dynamic Filters and Policers

[Dynamic Firewall Filters Overview | 225](#)

[Configuring Static Firewall Filters That Are Dynamically Applied | 229](#)

[Streamlining Processing of Chains of Static Filters | 237](#)

[Dynamically Attaching Static or Fast Update Filters to an Interface | 245](#)

[Configuring Filters That Are Created Dynamically | 249](#)

[Using Ascend Data Filters to Implement Firewalls Based on RADIUS Attributes | 303](#)

[Configuring Fast Update Filters to Provide More Efficient Processing Over Classic Static Filters | 323](#)

[Defending Against DoS and DDoS Attacks Using Unicast RPF and Fail Filters | 347](#)

[Improving Scaling and Performance of Filters on Static Subscriber Interfaces | 359](#)

[Configuring Dynamic Service Sets | 365](#)

[Configuring Rate-Limiting Premium and Non-Premium Traffic on an Interface Using Hierarchical Policers | 369](#)

[Monitoring and Managing Firewalls for Subscriber Access | 389](#)

Dynamic Firewall Filters Overview

IN THIS CHAPTER

- Understanding Dynamic Firewall Filters | 226
- Defining Dynamic Filter Processing Order | 227

Understanding Dynamic Firewall Filters

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 are static filters that are applied to an interface dynamically. They are compiled at commit time and then, when a service is activated, an interface-specific filter is created and attached to a logical interface. This dynamic application is performed by associating input or output filters with a dynamic profile. When triggered, a dynamic profile applies the filter to an interface. Because classic filters are static, they cannot contain subscriber-specific terms (also called rules).
- Parameterized filters allow you to implement customized filters for each subscriber session. In parameterized filters, you use variables to define a filter. When services are activated for a subscriber, actual values such as policing rates, destination addresses, or ports are substituted for the variables and are used to create filters.
- Ascend-Data-Filters allow you to create dynamic filters based on values received from the RADIUS server in the Ascend-Data-Filter attribute (RADIUS attribute 242). 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. When services are activated for a subscriber, a filter is created based on the values in the RADIUS attribute. You can also use Ascend-Data-Filters to create static filters by configuring the Ascend-Data-Filter attribute in a dynamic profile.
- Fast update filters are similar to classic filters. 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 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 | 229](#)

[Ascend-Data-Filter Policies for Subscriber Management Overview | 303](#)

[Parameterized Filters Overview | 249](#)

[Fast Update Filters Overview | 323](#)

[Dynamically Attaching Statically Created Filters for Any Interface Type | 246](#)

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 *Firewall Filters Overview* for information about firewall filters and how to create them.

2. Create a basic dynamic profile.

See *Configuring a Basic Dynamic Profile*.

3. Attach the filters to the dynamic profile.

See “[Dynamically Attaching Statically Created Filters for Any Interface Type](#)” on page 246, “[Dynamically Attaching Statically Created Filters for a Specific Interface Family Type](#)” on page 245, or “[Dynamically Attaching Filters Using RADIUS Variables](#)” on page 290.

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](#) | 229

Firewall Filters Overview

Configuring Static Firewall Filters That Are Dynamically Applied

IN THIS CHAPTER

- [Classic Filters Overview | 229](#)
- [Basic Classic Filter Syntax | 232](#)
- [Examples: Configuring Static Filters | 233](#)

Classic Filters Overview

IN THIS SECTION

- [Classic Filter Types | 230](#)
- [Classic Filter Components | 230](#)
- [Classic Filter Processing | 230](#)
- [Guidelines for Creating and Applying Classic Filters for Subscriber Interfaces | 231](#)

The dynamic firewall feature supports classic filters, which are static filters that are applied to an interface dynamically. They 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. This dynamic application is performed by associating input or output filters with a dynamic profile.

This overview covers:

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.

Guidelines for Creating and Applying Classic Filters for Subscriber Interfaces

Dynamic configuration of firewall filters is supported. 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:

- Dynamic application of only input and output filters is supported.
- 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

[Understanding Dynamic Firewall Filters | 226](#)

[Fast Update Filters Overview | 323](#)

[Dynamically Attaching Statically Created Filters for Any Interface Type | 246](#)

[Dynamically Attaching Statically Created Filters for a Specific Interface Family Type | 245](#)

[Dynamically Attaching Filters Using RADIUS Variables | 290](#)

[Verifying and Managing Firewall Filter Configuration | 389](#)

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 | 245](#)

[Understanding Dynamic Firewall Filters | 226](#)

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 {
                        192.51.100.10/24;
                    }
                }
                then {
                    count c1;
                    next term;
                }
            }
            term 2 {

```

```
        from {
            source-address {
                192.51.100.20/24;
            }
        }
        then count c2;
    }
    term 3 {
        then accept;
    }
}

filter dfwd1 {
    interface-specific;
    term 1 {
        from {
            address {
                192.51.100.10/24;
            }
        }
        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 {
        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 | 246](#)

[Dynamically Attaching Statically Created Filters for a Specific Interface Family Type | 245](#)

Streamlining Processing of Chains of Static Filters

IN THIS CHAPTER

- [Configuring Firewall Filter Bypass | 237](#)
- [Example: Bypassing Firewall Filters | 238](#)

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 227](#) 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.

```
[edit firewall family inet filter data term 1]
user@host# set then accept
```

RELATED DOCUMENTATION

[Classic Filters Overview | 229](#)

[Defining Dynamic Filter Processing Order | 227](#)

[Example: Bypassing Firewall Filters | 238](#)

Example: Bypassing Firewall Filters

IN THIS SECTION

- [Before You Begin | 239](#)
- [Filter Bypass Overview | 239](#)
- [Configuring Filter Bypass | 239](#)

This example describes how to configure multiple filters using the **service-filter-hit** match/action combination and contains the following sections:

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 227](#) for more information about dynamic filter processing and how to use the **precedence** statement.

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.

[Figure 5 on page 239](#) 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 5: Logical Flow Example for Filter Bypass Processing



Configuring Filter Bypass

IN THIS SECTION

- [Configuring the Voice Filter | 240](#)
- [Configuring the Video Filter | 240](#)
- [Configuring the Data Filter | 241](#)
- [Results | 241](#)

CLI Quick Configuration

To quickly configure this example:

```
[edit]
set firewall filter voice term T1 from address 203.0.113.11/32
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 203.0.113.100/32
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 5 on page 239](#):

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 203.0.113.11/32
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 5 on page 239](#):

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 203.0.113.100/32
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 5 on page 239](#):

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 {
```

```

term T1 {
  from {
    address {
      203.0.113.11/32;
    }
    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 {
        203.0.113.100/32;
      }
    }
    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 | 229](#)

[Defining Dynamic Filter Processing Order | 227](#)

[Configuring Firewall Filter Bypass | 237](#)

Dynamically Attaching Static or Fast Update Filters to an Interface

IN THIS CHAPTER

- Dynamically Attaching Statically Created Filters for a Specific Interface Family Type | 245
- Dynamically Attaching Statically Created Filters for Any Interface Type | 246

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 *Firewall Filters Overview* for information about classic firewall filters and how to create them. See [“Configuring Fast Update Filters” on page 328](#) for information about creating fast update filters.

2. Create a basic dynamic profile.

See *Configuring a Basic Dynamic Profile*.

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 | 229](#)

[Fast Update Filters Overview | 323](#)

[Dynamically Attaching Statically Created Filters for Any Interface Type | 246](#)

[Dynamically Attaching Filters Using RADIUS Variables | 290](#)

Using the junos-defaults Configuration Group

Firewall Filters Overview

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.

See *Firewall Filters Overview* for information about classic firewall filters and how to create them. See [“Configuring Fast Update Filters” on page 328](#) for information about creating fast update filters.

2. Create a basic dynamic profile.

See *Configuring a Basic Dynamic Profile*.

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](#) | 229

[Fast Update Filters Overview](#) | 323

[Dynamically Attaching Statically Created Filters for a Specific Interface Family Type](#) | 245

[Dynamically Attaching Filters Using RADIUS Variables](#) | 290

Using the junos-defaults Configuration Group

Firewall Filters Overview

Configuring Filters That Are Created Dynamically

IN THIS CHAPTER

- [Parameterized Filters Overview | 249](#)
- [Unique Identifiers for Firewall Variables | 250](#)
- [Configuring Unique Identifiers for Parameterized Filters | 253](#)
- [Sample Dynamic-Profile Configuration for Parameterized Filters | 254](#)
- [Dynamic Profile After UID Substitutions for Parameterized Filters | 256](#)
- [Multiple Parameterized Filters | 258](#)
- [Parameterized Filter Processing Overview | 259](#)
- [Parameterized Filters Configuration Considerations | 260](#)
- [Guidelines for Creating and Applying Parameterized Filters for Subscriber Interfaces | 262](#)
- [Parameterized Filter Match Conditions for IPv4 Traffic | 263](#)
- [Parameterized Filter Match Conditions for IPv6 Traffic | 270](#)
- [Parameterized Filter Nonterminating and Terminating Actions and Modifiers | 277](#)
- [Firewall Filter Match Conditions for Protocol-Independent Traffic in Dynamic Service Profiles | 284](#)
- [Firewall Filter Terminating and Nonterminating Actions for Protocol-Independent Traffic in Dynamic Service Profiles | 285](#)
- [Interface-Shared Filters Overview | 289](#)
- [Dynamically Attaching Filters Using RADIUS Variables | 290](#)
- [Example: Implementing a Filter for Households That Use ACI-Based VLANs | 292](#)
- [Example: Dynamic-Profile Parsing | 294](#)
- [Example: Firewall Dynamic Profile | 295](#)
- [Example: Configuring a Filter to Exclude DHCPv6 and ICMPv6 Control Traffic for LAC Subscriber | 296](#)

Parameterized Filters Overview

Parameterized filters allow you to implement customized filters for each subscriber session. In parameterized filters, you use variables called unique identifiers (UIDs) to define your filter. When services are activated for a subscriber, actual values are substituted for the variables and are used to create filters.

Parameterized filters are configured under a dynamic profile. You 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.

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. 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 to it.

RELATED DOCUMENTATION

[Parameterized Filters Configuration Considerations | 260](#)

[Parameterized Filter Processing Overview | 259](#)

[Unique Identifiers for Firewall Variables | 250](#)

[Sample Dynamic-Profile Configuration for Parameterized Filters | 254](#)

[Dynamic Profile After UID Substitutions for Parameterized Filters | 256](#)

[Example: Dynamic-Profile Parsing | 294](#)

[Parameterized Filter Nonterminating and Terminating Actions and Modifiers | 277](#)

[Parameterized Filter Match Conditions for IPv4 Traffic | 263](#)

[Parameterized Filter Match Conditions for IPv6 Traffic | 270](#)

[Understanding Dynamic Firewall Filters | 226](#)

Unique Identifiers for Firewall Variables

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 among multiple subscribers instead of creating a new filter 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;
    }
  }
}
```

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

[Configuring Unique Identifiers for Parameterized Filters | 253](#)

[Parameterized Filter Processing Overview | 259](#)

[Parameterized Filters Configuration Considerations | 260](#)

Configuring Unique Identifiers for Parameterized Filters

This topic discusses how to configure unique identifiers (UIDs) that can then be used in parameterized filters. The dynamic profile obtains and replaces data for these variables from an incoming client data packet.

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 UIDs.

If the value for the variable UID comes from RADIUS, configure the variable as a UID reference.

```
[edit dynamic-profiles Profile1]
user@host# set variable policer1 uid
user@host# set variable policer2 uid
user@host# set variable filter1 uid
user@host# set variable filter2 uid
user@host# set variables in-filter uid-reference
```

Example of UIDs that can be used in parameterized filters:

```
dynamic profile {
  Profile1 {
    variables {
      policer1 uid;
      filter1 uid;
      policer2 uid;
      filter2 uid;
      in-filter {
        uid-reference;
      }
    }
  }
}
```

[Unique Identifiers for Firewall Variables | 250](#)
[Parameterized Filters Overview | 249](#)
[Dynamic Variables Overview](#)
[Junos OS Predefined Variables](#)

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;
      }
    }
  }
  interfaces {
```



```

    }
    then {
        policer [$my-policer];
    }
}
term last {
    then {
        count drops;
        discard;
    }
}
}
}
}
}
}
}
}
}
}

```

RELATED DOCUMENTATION

[Dynamic Profile After UID Substitutions for Parameterized Filters | 256](#)

[Example: Dynamic-Profile Parsing | 294](#)

[Parameterized Filters Overview | 249](#)

[Parameterized Filter Processing Overview | 259](#)

Dynamic Profile After UID Substitutions for Parameterized Filters

In the following example, the client session is created on the ge-1/0/0.7 interface and this service is activated:

```
my-svc-prof(my-filt-lgw, 198.51.100.239/32, 0, 5m)
```

A dynamic profile is created by the process. The UIDs assigned by the process are based on the parameters being passed in as well as the sessions previously created.

```

dynamic-profile {
    [my-svc-prof] {
        interfaces {
            ge-1/0/0 {

```

```

    unit 7 {
        family inet {
            filter {
                input my-filt-1gw_UID1022;
            }
        }
    }
}

firewall {
    policer my-policer_UID1005 {
        if-exceeding {
            bandwidth-limit 5m;
            burst-size-limit 15000;
        }
        then discard;
    }
    family inet {
        filter my-filt-1gw_UID1022 {
            interface-specific;
            term t0 {
                from {
                    destination-address 198.51.100.239/32;
                }
                then {
                    policer my-policer_UID1005;
                }
            }
            term last {
                then {
                    count drops;
                    discard;
                }
            }
        }
    }
    filter my-filt-2gw_UID11018 {
        interface-specific;
        term t0 {
            from {
                destination-address {
                    198.51.100.239/32;
                    0;
                }
            }
        }
    }
}

```

```

        then {
            policer my-policer_UID1005;
        }
    }
    term last {
        then {
            count drops;
            discard;
        }
    }
}
}
}
}
}
}
}
}
}
}

```

RELATED DOCUMENTATION

[Sample Dynamic-Profile Configuration for Parameterized Filters | 254](#)

[Example: Dynamic-Profile Parsing | 294](#)

[Parameterized Filters Overview | 249](#)

[Parameterized Filter Processing Overview | 259](#)

Multiple Parameterized Filters

Differing filter match conditions can be achieved by allowing the filter that is being attached to be selected by the unique-identifier--reference capabilities of parameterized filters. If a variable number of terms or varying match conditions are needed, multiple filters are defined. When the service is activated, that activation will select the particular filter that should be applied in the stanza specifying the interface, unit, family and input/output filter:

```

interfaces {
    ge-1/0/0 {
        unit 7 {
            family inet {
                filter {
                    input my-filt-1gw-uid1022;
                }
            }
        }
    }
}

```

```

    }
  }
}

```

RELATED DOCUMENTATION

[Parameterized Filters Overview | 249](#)

[Parameterized Filters Configuration Considerations | 260](#)

Parameterized Filter Processing Overview

When creating a parameterized filter, you first define the family address type (**inet**, **inet6**, or **any**) 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 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](#) | 260

Parameterized Filters Configuration Considerations

IN THIS SECTION

- [Subscriber IP Address](#) | 261
- [Interaction with Static Configuration](#) | 261
- [Interface-Specific Dynamic Service Filters](#) | 261
- [Service Session Support](#) | 261
- [Filter Naming Conventions](#) | 261

Keep the following considerations in mind when configuring parameterized filters.

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 Dynamic Service Filters

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

- | |
|--|
| Understanding Dynamic Firewall Filters 226 |
| Verifying and Managing Firewall Filter Configuration 389 |
| Unique Identifiers for Firewall Variables 250 |
| Sample Dynamic-Profile Configuration for Parameterized Filters 254 |
| Example: Dynamic-Profile Parsing 294 |
| Parameterized Filter Processing Overview 259 |

Guidelines for Creating and Applying Parameterized Filters for Subscriber Interfaces

You can configure dynamic or static firewall filters. When you use statically configured firewall filters, you 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:

- Dynamic application of only input and output filters is supported.
- The filters must be interface-specific.
- You can create family-specific **any**, **inet**, and **inet6** filters.
- You can create interface-specific filters at the **unit** level that apply to any family type (**any**, **inet**, or **inet6**) configured on the interface.
- You can add or remove filters of different family types 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 259 |
| Parameterized Filters Configuration Considerations 260 |

Parameterized Filter Match Conditions for IPv4 Traffic

You can configure a parameterized filter with match conditions for Internet Protocol version 4 (IPv4) traffic (family inet).

NOTE: For MX Series routers with MPCs, you need to initialize certain new firewall filters by walking the corresponding SNMP MIB, for example, **show snmp mib walk name ascii**. This forces Junos to learn the filter counters and ensure that the filter statistics are displayed. This guidance applies to all enhanced mode firewall filters, filters with flexible conditions, and filters with certain terminating actions. See those topics, listed under Related Documentation, for details.

Table 23 on page 263 describes the *match-conditions* you can configure at the [edit firewall family inet filter filter-name term term-name from] hierarchy level.

Table 23: Parameterized Filter Match Conditions for IPv4 Traffic

Match Condition	Description
address address [except]	Match the IPv4 source or destination address field unless the except option is included. If the option is included, do not match the IPv4 source or destination address field.
destination-address address [except]	Match the IPv4 destination address field unless the except option is included. If the option is included, do not match the IPv4 destination address field. You cannot specify both the address and destination-address match conditions in the same term.

Table 23: Parameterized Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
destination-port <i>number</i>	<p>Match the UDP or TCP destination port field.</p> <p>You cannot specify both the port and destination-port match conditions in the same term.</p> <p>If you configure this match condition, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobileip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xdmcp (177).</p>
destination-port-except <i>number</i>	<p>Do not match the UDP or TCP destination port field. For details, see the destination-port match condition.</p>
destination-prefix-list <i>prefix-list-name</i> [<i>except</i>]	<p>Match destination prefixes in the specified list unless the except option is included. If the option is included, do not match the destination prefixes in the specified list.</p> <p>Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p>

Table 23: Parameterized Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
dscp number	<p>Match the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant 6 bits of this byte form the DSCP. For more information, see <i>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i>.</p> <p>Starting in Junos OS Release 13.3R7, support was added for filtering on Differentiated Services Code Point (DSCP) and forwarding class for Routing Engine sourced packets, including IS-IS packets encapsulated in generic routing encapsulation (GRE). Subsequently, when upgrading from a previous version of Junos OS where you have both a class of service (CoS) and firewall filter, and both include DSCP or forwarding class filter actions, the criteria in the firewall filter automatically takes precedence over the CoS settings. The same is true when creating new configurations; that is, where the same settings exist, the firewall filter takes precedence over the CoS, regardless of which was created first.</p> <p>You can specify a numeric value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • RFC 3246, <i>An Expedited Forwarding PHB (Per-Hop Behavior)</i>, defines one code point: ef (46). • RFC 2597, <i>Assured Forwarding PHB Group</i>, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: <ul style="list-style-type: none"> • af11 (10), af12 (12), af13 (14) • af21 (18), af22 (20), af23 (22) • af31 (26), af32 (28), af33 (30) • af41 (34), af42 (36), af43 (38)
dscp-except number	Do not match on the DSCP number. For more information, see the dscp match condition.
forwarding-class class	<p>Match the forwarding class of the packet.</p> <p>Specify assured-forwarding, best-effort, expedited-forwarding, or network-control.</p> <p>For information about forwarding classes and router-internal output queues, see <i>Understanding How Forwarding Classes Assign Classes to Output Queues</i>.</p>
forwarding-class-except class	Do not match the forwarding class of the packet. For details, see the forwarding-class match condition.

Table 23: Parameterized Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
icmp-code <i>number</i>	<p>Match the ICMP message code field.</p> <p>If you configure this match condition, we recommend that you also configure the protocol icmp match condition in the same term.</p> <p>If you configure this match condition, you must also configure the icmp-type <i>message-type</i> match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> • parameter-problem: ip-header-bad (0), required-option-missing (1) • redirect: redirect-for-host (1), redirect-for-network (0), redirect-for-tos-and-host (3), redirect-for-tos-and-net (2) • time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) • unreachable: communication-prohibited-by-filtering (13), destination-host-prohibited (10), destination-host-unknown (7), destination-network-prohibited (9), destination-network-unknown (6), fragmentation-needed (4), host-precedence-violation (14), host-unreachable (1), host-unreachable-for-TOS (12), network-unreachable (0), network-unreachable-for-TOS (11), port-unreachable (3), precedence-cutoff-in-effect (15), protocol-unreachable (2), source-host-isolated (8), source-route-failed (5)
icmp-code-except <i>message-code</i>	Do not match the ICMP message code field. For details, see the icmp-code match condition.
icmp-type <i>number</i>	<p>Match the ICMP message type field.</p> <p>If you configure this match condition, we recommend that you also configure the protocol icmp match condition in the same term.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): echo-reply (0), echo-request (8), info-reply (16), info-request (15), mask-request (17), mask-reply (18), parameter-problem (12), redirect (5), router-advertisement (9), router-solicit (10), source-quench (4), time-exceeded (11), timestamp (13), timestamp-reply (14), or unreachable (3).</p>
icmp-type-except <i>message-type</i>	Do not match the ICMP message type field. For details, see the icmp-type match condition.

Table 23: Parameterized Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
loss-priority level	<p>Match the packet loss priority (PLP) level.</p> <p>Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>For IP traffic on MX Series routers with Enhanced II Flexible PIC Concentrators (FPCs), you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement, see <i>Configuring and Applying Tricolor Marking Policies</i>. For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see <i>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i>.</p>
loss-priority-except level	Do not match the PLP level. For details, see the loss-priority match condition.
packet-length bytes	Match the length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.
packet-length-except bytes	Do not match the length of the received packet, in bytes. For details, see the packet-length match type.
port number	<p>Match the UDP or TCP source or destination port field.</p> <p>If you configure this match condition, you cannot configure the destination-port match condition or the source-port match condition in the same term.</p> <p>If you configure this match condition, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed under the destination-port match condition.</p>
port-except number	Do not match either the source or destination UDP or TCP port field. For details, see the port match condition.
precedence ip-precedence-value	<p>Match the IP precedence field.</p> <p>In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): critical-ecp (0xa0), flash (0x60), flash-override (0x80), immediate (0x40), internet-control (0xc0), net-control (0xe0), priority (0x20), or routine (0x00). You can specify precedence in hexadecimal, binary, or decimal form.</p>

Table 23: Parameterized Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
precedence-except ip-precedence-value	<p>Do not match the IP precedence field.</p> <p>In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): critical-ecp (0xa0), flash (0x60), flash-override (0x80), immediate (0x40), internet-control (0xc0), net-control (0xe0), priority (0x20), or routine (0x00). You can specify precedence in hexadecimal, binary, or decimal form.</p>
prefix-list prefix-list-name [except]	<p>Match the prefixes of the source or destination address fields to the prefixes in the specified list unless the except option is included. If the option is included, do not match the prefixes of the source or destination address fields to the prefixes in the specified list.</p> <p>The prefix list is defined at the [edit policy-options prefix-list prefix-list-name] hierarchy level.</p>
protocol number	<p>Match the IP protocol type field. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): ah (51), dstop (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmp6 (58), icmpv6 (58), igmp (2), ipip (4), ipv6 (41), ospf (89), pim (103), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).</p>
protocol-except number	<p>Do not match the IP protocol type field. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): ah (51), dstop (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmp6 (58), icmpv6 (58), igmp (2), ipip (4), ipv6 (41), ospf (89), pim (103), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).</p>
service-filter-hit	<p>Match a packet received from a filter where a service-filter-hit action was applied.</p>
source-address address [except]	<p>Match the IPv4 address of the source node sending the packet unless the except option is included. If the option is included, do not match the IPv4 address of the source node sending the packet.</p> <p>You cannot specify both the address and source-address match conditions in the same term.</p>
source-class class-names	<p>Match one or more specified source class names (sets of source prefixes grouped together and given a class name). For more information, see <i>Firewall Filter Match Conditions Based on Address Classes</i>.</p>
source-class-except class-names	<p>Do not match one or more specified source class names. For details, see the source-class match condition.</p>

Table 23: Parameterized Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
source-port number	<p>Match the UDP or TCP source port field.</p> <p>You cannot specify the port and source-port match conditions in the same term.</p> <p>If you configure this match condition for IPv4 traffic, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed with the destination-port number match condition.</p>
source-port-except number	Do not match the UDP or TCP source port field. For details, see the source-port match condition.
source-prefix-list name [except]	<p>Match source prefixes in the specified list unless the except option is included. If the option is included, do not match the source prefixes in the specified list.</p> <p>Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p>
ttl number	Match the IPv4 time-to-live number. Specify a TTL value or a range of TTL values. For number , you can specify one or more values from 0 through 255 . This match condition is supported only on M120, M320, MX Series, and T Series routers.
ttl-except number	Do not match on the IPv4 TTL number. For details, see the ttl match condition.

Release History Table

Release	Description
13.3R7	Starting in Junos OS Release 13.3R7, support was added for filtering on Differentiated Services Code Point (DSCP) and forwarding class for Routing Engine sourced packets, including IS-IS packets encapsulated in generic routing encapsulation (GRE).

RELATED DOCUMENTATION

[Parameterized Filters Overview | 249](#)
[Guidelines for Creating and Applying Parameterized Filters for Subscriber Interfaces | 262](#)
[Parameterized Filter Nonterminating and Terminating Actions and Modifiers | 277](#)

Parameterized Filter Match Conditions for IPv6 Traffic

You can configure a parameterized filter with match conditions for Internet Protocol version 6 (IPv6) traffic (family inet6).

NOTE: For MX Series routers with MPCs, you need to initialize certain new firewall filters by walking the corresponding SNMP MIB, for example, **show snmp mib walk name ascii**. This forces Junos to learn the filter counters and ensure that the filter statistics are displayed. This guidance applies to all enhanced mode firewall filters, filters with flexible conditions, and filters with certain terminating actions. See those topics, listed under Related Documentation, for details.

Table 24 on page 270 describes the match conditions you can configure at the [edit firewall family inet6 filter filter-name term term-name from] hierarchy level.

Table 24: Firewall Filter Match Conditions for IPv6 Traffic

Match Condition	Description
address address [except]	Match the IPv6 source or destination address field unless the except option is included. If the option is included, do not match the IPv6 source or destination address field.
destination-address address [except]	<p>Match the IPv6 destination address field unless the except option is included. If the option is included, do not match the IPv6 destination address field.</p> <p>You cannot specify both the address and destination-address match conditions in the same term.</p>

Table 24: Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
destination-port <i>number</i>	<p>Match the UDP or TCP destination port field.</p> <p>You cannot specify both the port and destination-port match conditions in the same term.</p> <p>If you configure this match condition, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobileip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xdmcp (177).</p>
destination-port-except <i>number</i>	Do not match the UDP or TCP destination port field. For details, see the destination-port match condition.
destination-prefix-list <i>prefix-list-name</i> [<i>except</i>]	<p>Match the IPv6 destination prefix to the specified list unless the except option is included. If the option is included, do not match the IPv6 destination prefix to the specified list.</p> <p>The prefix list is defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p>
forwarding-class <i>class</i>	<p>Match the forwarding class of the packet.</p> <p>Specify assured-forwarding, best-effort, expedited-forwarding, or network-control.</p> <p>For information about forwarding classes and router-internal output queues, see <i>Understanding How Forwarding Classes Assign Classes to Output Queues</i>.</p>
forwarding-class-except <i>class</i>	Do not match the forwarding class of the packet. For details, see the forwarding-class match condition.

Table 24: Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
icmp-code <i>message-code</i>	<p>Match the ICMP message code field.</p> <p>If you configure this match condition, we recommend that you also configure the next-header icmp or next-header icmp6 match condition in the same term.</p> <p>If you configure this match condition, you must also configure the icmp-type message-type match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem: ip6-header-bad (0), unrecognized-next-header (1), unrecognized-option (2) time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) destination-unreachable: administratively-prohibited (1), address-unreachable (3), no-route-to-destination (0), port-unreachable (4)
icmp-code-except <i>message-code</i>	Do not match the ICMP message code field. For details, see the icmp-code match condition.

Table 24: Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
icmp-type message-type	<p>Match the ICMP message type field.</p> <p>If you configure this match condition, we recommend that you also configure the next-header icmp or next-header icmp6 match condition in the same term.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): certificate-path-advertisement (149), certificate-path-solicitation (148), destination-unreachable (1), echo-reply (129), echo-request (128), home-agent-address-discovery-reply (145), home-agent-address-discovery-request (144), inverse-neighbor-discovery-advertisement (142), inverse-neighbor-discovery-solicitation (141), membership-query (130), membership-report (131), membership-termination (132), mobile-prefix-advertisement-reply (147), mobile-prefix-solicitation (146), neighbor-advertisement (136), neighbor-solicit (135), node-information-reply (140), node-information-request (139), packet-too-big (2), parameter-problem (4), private-experimentation-100 (100), private-experimentation-101 (101), private-experimentation-200 (200), private-experimentation-201 (201), redirect (137), router-advertisement (134), router-renumbering (138), router-solicit (133), or time-exceeded (3).</p> <p>For private-experimentation-201 (201), you can also specify a range of values within square brackets.</p>
icmp-type-except message-type	Do not match the ICMP message type field. For details, see the icmp-type match condition.
loss-priority level	<p>Match the packet loss priority (PLP) level.</p> <p>Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers and EX Series switches.</p> <p>For IP traffic on M320, MX Series, T Series routers and EX Series switches with Enhanced II Flexible PIC Concentrators (FPCs), you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement, see <i>Configuring and Applying Tricolor Marking Policers</i>. For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see <i>Understanding How Forwarding Classes Assign Classes to Output Queues</i>.</p>

Table 24: Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
loss-priority-except level	Do not match the PLP level. For details, see the loss-priority match condition.
next-header header-type	<p>Match the first 8-bit Next Header field in the packet. Support for the next-header firewall match condition is available in Junos OS Release 13.3R6 and later.</p> <p>For IPv6, we recommend that you use the payload-protocol term rather than the next-header term when configuring a firewall filter with match conditions. Although either can be used, payload-protocol provides the more reliable match condition because it uses the actual payload protocol to find a match, whereas next-header simply takes whatever appears in the first header following the IPv6 header, which may or may not be the actual protocol. In addition, if next-header is used with IPv6, the accelerated filter block lookup process is bypassed and the standard filter used instead.</p> <p>Match the first 8-bit Next Header field in the packet.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): ah (51), dstops (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmp6 (58), icmpv6 (58), igmp (2), ipip (4), ipv6 (41), mobility (135), no-next-header (59), ospf (89), pim (103), routing (43), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).</p> <p>NOTE: next-header icmp6 and next-header icmpv6 match conditions perform the same function. next-header icmp6 is the preferred option. next-header icmpv6 is hidden in the Junos OS CLI.</p>
next-header-except header-type	Do not match the 8-bit Next Header field that identifies the type of header between the IPv6 header and payload. For details, see the next-header match type.
packet-length bytes	Match the length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.
packet-length-except bytes	Do not match the length of the received packet, in bytes. For details, see the packet-length match type.

Table 24: Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
port number	<p>Match the UDP or TCP source or destination port field.</p> <p>If you configure this match condition, you cannot configure the destination-port match condition or the source-port match condition in the same term.</p> <p>If you configure this match condition, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed under the destination-port match condition.</p>
port-except number	Do not match the UDP or TCP source or destination port field. For details, see the port match condition.
prefix-list prefix-list-name [except]	<p>Match the prefixes of the source or destination address fields to the prefixes in the specified list unless the except option is included. If the option is included, do not match the prefixes of the source or destination address fields to the prefixes in the specified list.</p> <p>The prefix list is defined at the [edit policy-options prefix-list prefix-list-name] hierarchy level.</p>
service-filter-hit	Match a packet received from a filter where a service-filter-hit action was applied.
source-address address [except]	<p>Match the IPv6 address of the source node sending the packet unless the except option is included. If the option is included, do not match the IPv6 address of the source node sending the packet.</p> <p>You cannot specify both the address and source-address match conditions in the same term.</p>
source-class class-names	Match one or more specified source class names (sets of source prefixes grouped together and given a class name). For more information, see <i>Firewall Filter Match Conditions Based on Address Classes</i> .
source-class-except class-names	Do not match one or more specified source class names. For details, see the source-class match condition.

Table 24: Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
source-port number	<p>Match the UDP or TCP source port field.</p> <p>You cannot specify the port and source-port match conditions in the same term.</p> <p>If you configure this match condition, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed with the destination-port number match condition.</p>
source-port-except number	Do not match the UDP or TCP source port field. For details, see the source-port match condition.
source-prefix-list name [except]	<p>Match the IPv6 address prefix of the packet source field unless the except option is included. If the option is included, do not match the IPv6 address prefix of the packet source field.</p> <p>Specify a prefix list name defined at the [edit policy-options prefix-list prefix-list-name] hierarchy level.</p>
traffic-class number	<p>Match the 8-bit field that specifies the class-of-service (CoS) priority of the packet.</p> <p>This field was previously used as the type-of-service (ToS) field in IPv4.</p> <p>You can specify a numeric value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • RFC 3246, <i>An Expedited Forwarding PHB (Per-Hop Behavior)</i>, defines one code point: ef (46). • RFC 2597, <i>Assured Forwarding PHB Group</i>, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: <ul style="list-style-type: none"> • af11 (10), af12 (12), af13 (14) • af21 (18), af22 (20), af23 (22) • af31 (26), af32 (28), af33 (30) • af41 (34), af42 (36), af43 (38)
traffic-class-except number	Do not match the 8-bit field that specifies the CoS priority of the packet. For details, see the traffic-class match description.

NOTE: If you specify an IPv6 address in a match condition (the **address**, **destination-address**, or **source-address** match conditions), use the syntax for text representations described in RFC 4291, *IP Version 6 Addressing Architecture*. For more information about IPv6 addresses, see *IPv6 Overview and Supported IPv6 Standards*.

Release History Table

Release	Description
13.3R6	Support for the next-header firewall match condition is available in Junos OS Release 13.3R6 and later.

RELATED DOCUMENTATION

- Guidelines for Configuring Firewall Filters*
- Firewall Filter Terminating Actions*
- Firewall Filter Nonterminating Actions*
- Firewall Filter Match Conditions for IPv4 Traffic*
- [enhanced-mode | 842](#)
- Firewall Filter Flexible Match Conditions*

Parameterized Filter Nonterminating and Terminating Actions and Modifiers

The nonterminating and terminating actions and modifiers for parameterized filters are a subset of those available for static firewall filters.

NOTE: You cannot configure the **next term nonterminating** action with a *terminating* action in the same filter term. However, you can configure the **next term** action with another *nonterminating* action in the same filter term.

Nonterminating actions carry with them an implicit accept action. In this context, *nonterminating* means that other actions can follow these actions, whereas no other actions can follow a *terminating* action.

Table 25 on page 278 describes the nonterminating actions and modifiers you can configure for a parameterized filter term.

Table 25: Nonterminating Actions for Parameterized Filters

Nonterminating Action	Description	Protocol Families
count <i>counter-name</i>	Count the packet in the named counter.	<ul style="list-style-type: none"> • family any • family inet • family inet6
dscp value	<p>Set the IPv4 Differentiated Services code point (DSCP) bit. You can specify a numerical value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>The default DSCP value is best effort, that is, be or 0.</p> <p>You can also specify one of the following text synonyms:</p> <ul style="list-style-type: none"> • af11—Assured forwarding class 1, low drop precedence • af12—Assured forwarding class 1, medium drop precedence • af13—Assured forwarding class 1, high drop precedence • af21—Assured forwarding class 2, low drop precedence • af22—Assured forwarding class 2, medium drop precedence • af23—Assured forwarding class 2, high drop precedence • af31—Assured forwarding class 3, low drop precedence • af32—Assured forwarding class 3, medium drop precedence • af33—Assured forwarding class 3, high drop precedence • af41—Assured forwarding class 4, low drop precedence • af42—Assured forwarding class 4, medium drop precedence • af43—Assured forwarding class 4, high drop precedence • be—Best effort • cs0—Class selector 0 • cs1—Class selector 1 • cs2—Class selector 2 • cs3—Class selector 3 • cs4—Class selector 4 • cs5—Class selector 5 • cs6—Class selector 6 • cs7—Class selector 7 • ef—Expedited forwarding 	family inet

Table 25: Nonterminating Actions for Parameterized Filters (*continued*)

Nonterminating Action	Description	Protocol Families
forwarding-class <i>class-name</i>	<p>Classify the packet to the named forwarding class:</p> <ul style="list-style-type: none"> • assured-forwarding • best-effort • expedited-forwarding • network-control 	<ul style="list-style-type: none"> • family any • family inet • family inet6
hierarchical-policer	Police the packet using the specified hierarchical policer.	<ul style="list-style-type: none"> • family any • family inet • family inet6
log	<p>Log the packet header information in a buffer within the Packet Forwarding Engine. You can access this information by issuing the show firewall log command at the CLI.</p> <p>NOTE: The Layer 2 (L2) families log action is available only for MX Series routers with MPCs (MPC mode if the router has only MPCs, or mix mode if it has MPCs and DCPs). For MX Series routers with DPCs, the log action for L2 families is ignored if configured.</p>	<ul style="list-style-type: none"> • family inet • family inet6
loss-priority (high medium-high medium-low low)	<p>Set the packet loss priority (PLP) level.</p> <p>You cannot also configure the three-color-policer nonterminating action for the same firewall filter term. These two nonterminating actions are mutually exclusive.</p> <p>For IP traffic on MX Series routers with Enhanced II Flexible PIC Concentrators (FPCs), you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement and using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see <i>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i>.</p>	<ul style="list-style-type: none"> • family any • family inet • family inet6
next	Proceed to the next filter term.	<ul style="list-style-type: none"> • family any • family inet • family inet6

Table 25: Nonterminating Actions for Parameterized Filters (*continued*)

Nonterminating Action	Description	Protocol Families
next-ip <i>ip-address</i> <routing-instance <i>routing-instance></i>	<p>(MX Series) Direct packets to the specified destination IPv4 address. You can optionally specify a routing instance for the address. In the following example, the variables \$IP-address and \$RT-name would be defined in [edit dynamic-profiles <i>service-profile-name</i> variables]:</p> <pre>[edit dynamic-profiles <i>service-profile-name</i> firewall family inet filter \$nextip] user@host# set term t1 then next-ip \$IP-address routing-instance \$RT-name</pre> <p>Supported starting in Junos OS Release 18.2R1.</p>	family inet
next-ip6 <i>ipv6-address</i> <routing-instance <i>routing-instance></i>	<p>(MX Series) Direct packets to the specified destination IPv6 address. You can optionally specify a routing instance for the address. In the following example, the variables \$IPv6-address and \$RT-name would be defined in [edit dynamic-profiles <i>service-profile-name</i> variables]</p> <pre>[edit dynamic-profiles <i>service-profile-name</i> firewall family inet filter \$nextip6] user@host# set term t1 then next-ip6 \$IPv6-address routing-instance \$RT-name</pre> <p>Supported starting in Junos OS Release 18.2R1.</p>	family inet6
policer <i>policer-name</i>	Name of policer to use to rate-limit traffic.	<ul style="list-style-type: none"> • family any • family inet • family inet6
port-mirror <i>instance-name</i>	<p>Port-mirror the packet based on the specified family.</p> <p>We recommend that you do not use both the next-hop-group and the port-mirror actions in the same firewall filter.</p>	<ul style="list-style-type: none"> • family any • family inet • family inet6
port-mirror-instance <i>instance-name</i>	<p>Port-mirror a packet for an instance. This action is supported only on the MX Series routers.</p> <p>We recommend that you do not use both the next-hop-group and the port-mirror-instance actions in the same firewall filter.</p>	<ul style="list-style-type: none"> • family any • family inet • family inet6
routing-instance <i>routing-instance-name</i>	Direct packets to the specified routing instance.	<ul style="list-style-type: none"> • family inet • family inet6

Table 25: Nonterminating Actions for Parameterized Filters (*continued*)

Nonterminating Action	Description	Protocol Families
sample	<p>Sample the packet.</p> <p>NOTE: Junos OS does not sample packets originating from the router. If you configure a filter and apply it to the output side of an interface, then only the transit packets going through that interface are sampled. Packets that are sent from the Routing Engine to the Packet Forwarding Engine are not sampled.</p>	<ul style="list-style-type: none"> • family inet • family inet6
service-accounting	<p>Use the inline counting mechanism when capturing subscriber per-service statistics.</p> <p>Count the packet for service accounting. The count is applied to a specific named counter (_junos-dyn-service-counter) that RADIUS can obtain.</p> <p>The service-accounting and service-accounting-deferred keywords are mutually exclusive, both per-term and per-filter.</p>	<ul style="list-style-type: none"> • family any • family inet • family inet6
service-accounting-deferred	<p>Use the deferred counting mechanism when capturing subscriber per-service statistics. The count is applied to a specific named counter (_junos-dyn-service-counter) that RADIUS can obtain.</p> <p>The service-accounting and service-accounting-deferred keywords are mutually exclusive, both per-term and per-filter.</p>	<ul style="list-style-type: none"> • family any • family inet • family inet6
service-filter-hit	<p>(Only if the service-filter-hit flag is marked by a previous filter in the current type of chained filters) Direct the packet to the next type of filters.</p> <p>Indicate to subsequent filters in the chain that the packet was already processed. This action, coupled with the service-filter-hit match condition in receiving filters, helps to streamline filter processing.</p>	<ul style="list-style-type: none"> • family any • family inet • family inet6
three-color-policer (single-rate two-rate) policer-name	<p>Police the packet using the specified single-rate or two-rate three-color-policer.</p> <p>NOTE: You cannot also configure the loss-priority action for the same firewall filter term. These two actions are mutually exclusive.</p>	<ul style="list-style-type: none"> • family any • family inet • family inet6

Table 25: Nonterminating Actions for Parameterized Filters (*continued*)

Nonterminating Action	Description	Protocol Families
traffic-class value	<p>Specify the traffic-class code point. You can specify a numerical value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>The default traffic-class value is best effort, that is, be or 0.</p> <p>In place of the numeric value, you can specify one of the following text synonyms:</p> <ul style="list-style-type: none"> • af11—Assured forwarding class 1, low drop precedence • af12—Assured forwarding class 1, medium drop precedence • af13—Assured forwarding class 1, high drop precedence • af21—Assured forwarding class 2, low drop precedence • af22—Assured forwarding class 2, medium drop precedence • af23—Assured forwarding class 2, high drop precedence • af31—Assured forwarding class 3, low drop precedence • af32—Assured forwarding class 3, medium drop precedence • af33—Assured forwarding class 3, high drop precedence • af41—Assured forwarding class 4, low drop precedence • af42—Assured forwarding class 4, medium drop precedence • af43—Assured forwarding class 4, high drop precedence • be—Best effort • cs0—Class selector 0 • cs1—Class selector 1 • cs2—Class selector 2 • cs3—Class selector 3 • cs4—Class selector 4 • cs5—Class selector 5 • cs6—Class selector 6 • cs7—Class selector 7 • ef—Expedited forwarding 	family inet6

[Table 26 on page 283](#) describes the terminating actions and modifiers you can configure for a parameterized filter term.

Table 26: Terminating Actions for Parameterized Filters

Terminating Action	Description	Protocol Families
accept	Accept the packet.	<ul style="list-style-type: none"> • family any • family inet • family inet6
discard	Discard a packet silently, without sending an Internet Control Message Protocol (ICMP) message. Discarded packets are available for logging and sampling.	<ul style="list-style-type: none"> • family any • family inet • family inet6
reject <i>message-type</i>	<p>Reject the packet and return an ICMPv4 or ICMPv6 message:</p> <ul style="list-style-type: none"> • If no <i>message-type</i> is specified, a destination unreachable message is returned by default. • If tcp-reset is specified as the <i>message-type</i>, tcp-reset is returned only if the packet is a TCP packet. Otherwise, the administratively-prohibited message, which has a value of 13, is returned. • If any other <i>message-type</i> is specified, that message is returned. <p>NOTE: Rejected packets can be sampled or logged if you configure the sample or syslog action.</p> <p>The <i>message-type</i> can be one of the following values: address-unreachable, administratively-prohibited, bad-host-tos, bad-network-tos, beyond-scope, fragmentation-needed, host-prohibited, host-unknown, host-unreachable, network-prohibited, network-unknown, network-unreachable, no-route, port-unreachable, precedence-cutoff, precedence-violation, protocol-unreachable, source-host-isolated, source-route-failed, or tcp-reset.</p>	<ul style="list-style-type: none"> • family inet • family inet6

RELATED DOCUMENTATION

[Parameterized Filters Overview | 249](#)
[Guidelines for Creating and Applying Parameterized Filters for Subscriber Interfaces | 262](#)
[Parameterized Filter Match Conditions for IPv4 Traffic | 263](#)
[Parameterized Filter Match Conditions for IPv6 Traffic | 270](#)
[Understanding Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address](#)

Firewall Filter Match Conditions for Protocol-Independent Traffic in Dynamic Service Profiles

You configure firewall filter match conditions to determine which packets are filtered. Starting in Junos OS Release 16.1, you can configure match conditions that are supported for protocol-independent traffic—that is, configured under **family any**—for filters in dynamic service profiles. [Table 27 on page 284](#) describes these match conditions.

NOTE: Protocol-independent firewall filters in dynamic service profiles are supported only on MX Series routers with MPCs.

Table 27: Firewall Filter Match Conditions for Protocol-Independent Traffic in Dynamic Service Profiles

Match Condition	Description
forwarding-class class	<p>Match the forwarding class of the packet.</p> <p>Specify assured-forwarding, best-effort, expedited-forwarding, or network-control.</p> <p>For information about forwarding classes and router-internal output queues, see <i>Understanding How Forwarding Classes Assign Classes to Output Queues</i>.</p>
forwarding-class-except class	Do not match on the forwarding class. For details, see the forwarding-class match condition.
loss-priority level	<p>Match the packet loss priority (PLP) level.</p> <p>Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>For information about the tri-color statement, see <i>Configuring and Applying Tricolor Marking Policers</i>. For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see <i>Understanding How Forwarding Classes Assign Classes to Output Queues</i>.</p>
loss-priority-except level	Do not match the PLP level. For details, see the loss-priority match condition.
packet-length bytes	Match the length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.
packet-length-except bytes	Do not match on the received packet length, in bytes. For details, see the packet-length match type.

Table 27: Firewall Filter Match Conditions for Protocol-Independent Traffic in Dynamic Service Profiles (continued)

Match Condition	Description
service-filter-hit	<p>(Only if the service-filter-hit flag is marked by a previous filter in the current type of chained filters) Direct the packet to the next type of filters.</p> <p>Indicate to subsequent filters in the chain that the packet was already processed. This match option, coupled with the service-filter-hit nonterminating action, helps to streamline filter processing.</p>

Release History Table

Release	Description
16.1	Starting in Junos OS Release 16.1, you can configure match conditions that are supported for protocol-independent traffic—that is, configured under family any —for filters in dynamic service profiles.

RELATED DOCUMENTATION

<i>Guidelines for Configuring Firewall Filters</i>
Firewall Filter Terminating and Nonterminating Actions for Protocol-Independent Traffic in Dynamic Service Profiles 285
<i>Firewall Filter Match Conditions for IPv4 Traffic</i>
<i>Firewall Filter Match Conditions for IPv6 Traffic</i>

Firewall Filter Terminating and Nonterminating Actions for Protocol-Independent Traffic in Dynamic Service Profiles

Firewall filters in dynamic service profiles support a set of terminating actions that halt all evaluation of a firewall filter for a specific packet. The router performs the specified action, and no additional terms are examined. [Table 28 on page 286](#) describes the terminating actions conditions that are supported for protocol-independent traffic—that is, configured under **family any**—for filters in dynamic service profiles.

NOTE: You cannot configure the **next** action with a *terminating* action in the same filter term. However, you can configure the **next** action with another *nonterminating* action in the same filter term.

NOTE: Protocol-independent firewall filters in dynamic service profiles are supported only on MX Series routers with MPCs.

Table 28: Terminating Actions for Firewall Filters for Protocol-Independent Traffic in Dynamic Service Profiles

Terminating Action	Description
accept	Accept the packet.
discard	Discard a packet silently, without sending an Internet Control Message Protocol (ICMP) message. Discarded packets are available for logging and sampling.

Firewall filters in dynamic service profiles also support a set of nonterminating actions that are performed for a specific packet before the packet is passed to any subsequent actions in the term. [Table 28 on page 286](#) describes the terminating actions conditions that are supported for protocol-independent traffic—that is, configured under **family any**—for filters in dynamic service profiles.

Table 29: Nonterminating Actions for Firewall Filters for Protocol-Independent Traffic in Dynamic Service Profiles

Nonterminating Action	Description
count <i>counter-name</i>	Count the packet in the named counter.
force-premium	<p>By default, a hierarchical policer processes the traffic it receives according to the traffic's forwarding class. Premium, expedited-forwarding traffic, has priority for bandwidth over aggregate, best-effort traffic. The force-premium filter ensures that traffic matching the term is treated as premium traffic by a subsequent hierarchical policer, regardless of its forwarding class. This traffic is given preference over any aggregate traffic received by that policer.</p> <p>NOTE: The force-premium filter option is supported only on MPCs.</p>

Table 29: Nonterminating Actions for Firewall Filters for Protocol-Independent Traffic in Dynamic Service Profiles (continued)

Nonterminating Action	Description
forwarding-class <i>class-name</i>	<p>Classify the packet to the named forwarding class:</p> <ul style="list-style-type: none"> • <i>forwarding-class-name</i> • assured-forwarding • best-effort • expedited-forwarding • network-control <p>NOTE: This action is supported on ingress only.</p>
hierarchical-policer	Police the packet using the specified hierarchical policer.
loss-priority (high medium-high medium-low low)	<p>Set the packet loss priority (PLP) level.</p> <p>You cannot also configure the three-color-policer nonterminating action for the same firewall filter term. These two nonterminating actions are mutually exclusive.</p> <p>You must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can configure only the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement, see <i>Configuring and Applying Tricolor Marking Policers</i>. For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see <i>Understanding How Forwarding Classes Assign Classes to Output Queues</i>.</p> <p>For information about the tri-color statement and using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see <i>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i>.</p> <p>NOTE: This action is supported on ingress only.</p>
next	Proceed to the next filter term.
policer <i>policer-name</i>	Name of policer to use to rate-limit traffic.
port-mirror <i>instance-name</i>	<p>Port-mirror the packet based on the specified family.</p> <p>NOTE: This action is supported on ingress only.</p>

Table 29: Nonterminating Actions for Firewall Filters for Protocol-Independent Traffic in Dynamic Service Profiles (continued)

Nonterminating Action	Description
service-accounting	<p>Use the inline counting mechanism when capturing subscriber per-service statistics.</p> <p>Count the packet for service accounting. The count is applied to a specific named counter (_junos-dyn-service-counter) that RADIUS can obtain.</p> <p>The service-accounting and service-accounting-deferred keywords are mutually exclusive, both per-term and per-filter.</p> <p>NOTE: This action is not supported on T4000 Type 5 FPCs and PTX Series Packet Transport Routers.</p>
service-accounting-deferred	<p>Use the deferred counting mechanism when capturing subscriber per-service statistics. The count is applied to a specific named counter (_junos-dyn-service-counter) that RADIUS can obtain.</p> <p>The service-accounting and service-accounting-deferred keywords are mutually exclusive, both per-term and per-filter.</p> <p>NOTE: This action is not supported on T4000 Type 5 FPCs and PTX Series Packet Transport Routers.</p>
service-filter-hit	<p>(Only if the service-filter-hit flag is marked by a previous filter in the current type of chained filters) Direct the packet to the next type of filters.</p> <p>Indicate to subsequent filters in the chain that the packet was already processed. This action, coupled with the service-filter-hit match condition in receiving filters, helps to streamline filter processing.</p> <p>NOTE: This action is not supported on T4000 Type 5 FPCs and PTX Series Packet Transport Routers.</p>
three-color-policer (single-rate two-rate) <i>policer-name</i>	<p>Police the packet using the specified single-rate or two-rate three-color-policer.</p> <p>NOTE: You cannot also configure the loss-priority action for the same firewall filter term. These two actions are mutually exclusive.</p>

RELATED DOCUMENTATION

Guidelines for Configuring Firewall Filters

Firewall Filter Nonterminating Actions

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 be taken from either the **\$junos-interface-set-name** variable or the **\$junos-svlan-interface-set-name** variable, where the values of the variables are provided by the related client session or service session. For example, if the **\$junos-interface-set-name** variable is defined as the shared name, the same interface-shared filter instance is attached to all logical interfaces that belong to the interface set defined by the variable of that session. Similarly, if **\$junos-svlan-interface-set-name** is defined for the shared name, all logical interfaces that belong to the VLAN interfaces set defined by the session's variable share the same interface-shared instance.

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: Implementing a Filter for Households That Use ACI-Based VLANs | 292](#)

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
26-191	Input-Interface-Filter	\$junos-input-interface-filter
26-192	Output-Interface-Filter	\$junos-output-interface-filter

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
```

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
```

To dynamically attach input and output filters to any interface independent of protocol using RADIUS:

1. Specify the dynamic profile you want to attach, the interface, and the logical unit number.

```
[edit]
user@host# edit dynamic-profiles myProfile interfaces ge-1/1/1 unit 1
```

2. Specify the input filter variable that applies to all families configured for the logical interface.

```
[edit dynamic-profiles myProfile interfaces ge-1/1/1 unit 1]
user@host# set filter input $junos-input-interface-filter
```

3. Specify the output filter variable that applies to all families configured for the logical interface.

```
[edit dynamic-profiles myProfile interfaces ge-1/1/1 unit 1]
user@host# set filter output $junos-output-interface-filter
```

RELATED DOCUMENTATION

[Classic Filters Overview | 229](#)

[Dynamically Attaching Statically Created Filters for Any Interface Type | 246](#)

[Dynamically Attaching Statically Created Filters for a Specific Interface Family Type | 245](#)

[Junos OS Predefined Variables](#)

[Using the junos-defaults Configuration Group](#)

Example: Implementing a Filter for Households That Use ACI-Based VLANs

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, give a higher precedence to any other filters that are attached to the same interfaces.

Filter with interface-set match cannot be used on dynamic interface—dynamic interface-set match is not supported. The shared-name of an interface-shared filter can now be populated from the **\$junos-svlan-interface-set-name** variable. This means interface-shared filter can also be attached to dynamic SVLAN interface-set, before which the shared-name could only be taken from the **\$junos-interface-set-name** variable.

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 client-profile
```

2. Specify the interfaces and the unit.

```
[edit dynamic-profiles client-profile]
user@host# edit interfaces demux0 unit $junos-interface-unit
```

3. Specify the family.

```
[edit dynamic-profiles client-profile interfaces demux0 unit "$junos-interface-unit"]
user@host# edit family inet
```

4. Specify the input filter and the filter terms for the interface unit.

```
[edit dynamic-profiles client-profile interfaces demux0 unit "$junos-interface-unit" family inet]
```

```
user@host# edit input $junos-input-filter shared-name $junos-interface-set-name precedence
precedence-number
```

5. Specify the output filter and the filter terms for the interface unit.

```
[edit dynamic-profiles client-profile interfaces demux0 unit "$junos-interface-unit" family inet]
user@host# edit input $junos-output-filter shared-name $junos-interface-set-name precedence
precedence-number
```

6. Specify that you want to configure a firewall, and specify the family.

```
[edit dynamic-profiles client-profile]
user@host# edit firewall family inet
```

7. Specify the filter.

```
[edit dynamic-profiles client-profile firewall family inet]
user@host# edit filter uid
```

8. Specify that the filter is an interface-shared filter.

```
[edit dynamic-profiles client-profile firewall family inet filter uid]
user@host# set interface-shared
```

```
[edit]
dynamic-profile {
  client-profile {
    interfaces {
      demux0 {
        unit $junos-interface-unit {
          family inet {
            filter {
              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 | 245](#)

[Dynamically Attaching Filters Using RADIUS Variables | 290](#)

Firewall Filters Overview

Example: Dynamic-Profile Parsing

The following example shows the basic dynamic-profile parsing steps for parameterized filters.

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 **198.51.100.239** 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.
 - 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	254
Dynamic Profile After UID Substitutions for Parameterized Filters	256

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
<code>\$junos-input-filter</code>	Ingress-Policy-Name	26-10
<code>\$junos-output-filter</code>	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

[Understanding Dynamic Firewall Filters](#) | 226

Example: Configuring a Filter to Exclude DHCPv6 and ICMPv6 Control Traffic for LAC Subscriber

IN THIS SECTION

- [Requirements](#) | 297
- [Overview](#) | 297
- [Configuration](#) | 297

This example shows how to configure a standard stateless firewall filter that excludes DHCPv6 and ICMPv6 control packets from being considered for idle-timeout detection for tunneled subscribers at the LAC.

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

Subscriber access on a LAC can be limited by configuring an idle timeout period that specifies the maximum period of time a subscriber can remain idle after the subscriber session is established. The LAC monitors the subscriber's upstream and downstream data traffic to determine whether the subscriber is inactive. Based on the session accounting statistics, the subscriber is not considered idle as long as data traffic is detected in either direction. When no traffic is detected for the duration of the idle time out, the subscriber is logged out gracefully similarly to a RADIUS-initiated disconnect or a CLI-initiated logout.

However, after a tunnel is established for L2TP subscribers, all packets through the tunnel at the LAC are treated as data packets. Consequently, the accounting statistics for the session are inaccurate and the subscriber is not considered to be idle as long as DHCPv6 and ICMPv6 control packets are being sent.

Starting in Junos OS Release 17.2R1, you can define a firewall filter for the **inet6** family with terms to match on these control packets. Include the use the **exclude-accounting** terminating action in the filter terms to drop these control packets.

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.

```
set access profile v6-exclude-idle session-options client-idle-timeout 10
set access profile v6-exclude-idle session-options client-idle-timeout-ingress-only
edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER
set interface-specific
set term EXCLUDE-ACCT-DHCP-INET6 from next-header udp
set term EXCLUDE-ACCT-DHCP-INET6 from source-port 546
set term EXCLUDE-ACCT-DHCP-INET6 from source-port 547
set term EXCLUDE-ACCT-DHCP-INET6 from destination-port 546
set term EXCLUDE-ACCT-DHCP-INET6 from destination-port 547
set term EXCLUDE-ACCT-DHCP-INET6 then count exclude-acct-dhcpv6
set term EXCLUDE-ACCT-DHCP-INET6 then exclude-accounting
set term EXCLUDE-ACCT-ICMP6 from next-header icmp6
set term EXCLUDE-ACCT-ICMP6 from icmp-type router-solicit
set term EXCLUDE-ACCT-ICMP6 from icmp-type neighbor-solicit
set term EXCLUDE-ACCT-ICMP6 from icmp-type neighbor-advertisement
```

```

set term EXCLUDE-ACCT-ICMP6 then count exclude-acct-icmpv6
set term EXCLUDE-ACCT-ICMP6 then exclude-accounting
set term default then accept
top
edit dynamic-profiles pppoe-dynamic-profile interfaces pp0 unit "$junos-interface-unit"
set family inet6 filter input EXCLUDE-ACCT-INET6-FILTER
set family inet6 filter output EXCLUDE-ACCT-INET6-FILTER
set actual-transit-statistics

```

Configure the Filter

Step-by-Step Procedure

The following example requires that you 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 filter:

1. Set the idle timeout for subscriber sessions..

```

[edit access profile v6-exclude-idle]
user@host# set session-options client-idle-timeout 10

```

2. Specify the idle timeout applies only to ingress traffic.

```

[edit access profile v6-exclude-idle]
user@host# set session-options client-idle-timeout-ingress-only

```

3. Define the firewall filter term that excludes the DHCPv6 control packets from accounting statistics.
 - a. Specify a match on packets with the first Next Header field set to UDP (17).

```

[edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER]
user@host# set term EXCLUDE-ACCT-DHCP-INET6 from next-header udp

```

- b. Specify a match on packets with a source port of 546 or 547 (DHCPv6).

```

[edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER]
user@host# set term EXCLUDE-ACCT-DHCP-INET6 from source-port 546
user@host# set term EXCLUDE-ACCT-DHCP-INET6 from source-port 547

```

- c. Specify a match on packets with a DHCP destination port of 546 or 547 (DHCPv6).

```
[edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER]
user@host# set term EXCLUDE-ACCT-DHCP-INET6 from destination-port 546
user@host# set term EXCLUDE-ACCT-DHCP-INET6 from destination-port 547
```

- d. Count the matched DHCPv6 packets.

```
[edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER]
user@host# set term EXCLUDE-ACCT-DHCP-INET6 then count exclude-acct-dhcpv6
```

- e. Exclude the matched DHCPv6 packets from accounting statistics.

```
[edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER]
user@host# set term EXCLUDE-ACCT-DHCP-INET6 then exclude-accounting
```

4. Define the firewall filter term that excludes the ICMPv6 control packets from accounting statistics.

- a. Specify a match on packets with the first Next Header field set to ICMPv6 (58).

```
[edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER]
user@host# set term EXCLUDE-ACCT-ICMP6 from next-header icmp6
```

- b. Specify a match on packets with an ICMPv6 message type.

```
[edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER]
user@host# set term EXCLUDE-ACCT-ICMP6 from icmp-type router-solicit
user@host# set term EXCLUDE-ACCT-ICMP6 from icmp-type neighbor-solicit
user@host# set term EXCLUDE-ACCT-ICMP6 from icmp-type neighbor-advertisement
```

- c. Count the matched ICMPv6 packets.

```
[edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER]
user@host# set term EXCLUDE-ACCT-ICMP6 then count exclude-acct-icmpv6
```

- d. Exclude the matched ICMPv6 packets from accounting statistics.

```
[edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER]
user@host# set term EXCLUDE-ACCT-DHCP-INET6 then exclude-accounting
```

5. Define the default filter term to accept all other packets.

```
[edit firewall family inet6 filter EXCLUDE-ACCT-INET6-FILTER]
user@host# set term default then accept
```

6. Configure the dynamic profile to apply the filter to input and output interfaces for the **inet6** family.

```
[edit dynamic-profiles pppoe-dynamic-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set family inet6 filter input EXCLUDE-ACCT-INET6-FILTER
user@host# set family inet6 filter output EXCLUDE-ACCT-INET6-FILTER
```

7. Enable subscriber management accurate accounting.

```
[edit dynamic-profiles pppoe-dynamic-profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set actual-transit-statistics
```

Results

From configuration mode, confirm your configuration by entering the **show access**, **show firewall**, and **show dynamic-profiles** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show access
profile v6-exclude-idle {
  session-options {
    client-idle-timeout 10;
    client-idle-timeout-ingress-only;
  }
}
```

```
user@host# show firewall
family inet6 {
  filter EXCLUDE-ACCT-INET6-FILTER {
    interface-specific;
    term EXCLUDE-ACCT-DHCP-INET6 {
      from {
        next-header udp;
        source-port [ 546 547 ];
        destination-port [ 546 547 ];
      }
      then {
        count exclude-acct-dhcpv6;
        exclude-accounting
      }
    }
  }
  term EXCLUDE-ACCT-ICMP6 {
```

```

    from {
        next-header icmp6;
        icmp-type [ router-solicit neighbor-solicit neighbor-advertisement ]
    }
    then {
        count exclude-acct-icmpv6;
        exclude-accounting;
    }
}
term default {
    then accept;
}
}
}

```

```

user@host# show dynamic-profiles
pppoe-dynamic-profile {
    interfaces {
        pp0 {
            unit "$junos-interface-unit" {
                actual-transit-statistics;
                family inet6 {
                    filter {
                        input EXCLUDE-ACCT-INET6-FILTER;
                        output EXCLUDE-ACCT-INET6-FILTER;
                    }
                }
            }
        }
    }
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

RELATED DOCUMENTATION

[Classic Filters Overview | 229](#)

[Dynamically Attaching Statically Created Filters for a Specific Interface Family Type | 245](#)

[Understanding How to Use Standard Firewall Filters](#)

[Firewall Filter Terminating Actions](#)

[Firewall Filter Match Conditions for IPv6 Traffic](#)

Using Ascend Data Filters to Implement Firewalls Based on RADIUS Attributes

IN THIS CHAPTER

- [Ascend-Data-Filter Policies for Subscriber Management Overview | 303](#)
- [Ascend-Data-Filter Attribute Fields | 305](#)
- [Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions | 308](#)
- [Example: Configuring Dynamic Ascend-Data-Filter Support for Subscriber Access | 310](#)
- [Example: Configuring Static Ascend-Data-Filter Support for Subscriber Access | 315](#)
- [Verifying and Managing Dynamic Ascend-Data-Filter Policy Configuration | 320](#)

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.

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**, ..., **tn**. 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 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 | 308](#)
- [Ascend-Data-Filter Attribute Fields | 305](#)

Ascend-Data-Filter Attribute Fields

[Table 30 on page 305](#) 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 30: Ascend-Data-Filter Attribute Fields

Action or Classifier	Format	Value	Junos OS Filter Function
Type	1 byte	<ul style="list-style-type: none"> • 1 = IPv4 • 3 = IPv6 	
Filter or forward	1 byte	<ul style="list-style-type: none"> • 0 = filter • 1 = forward 	<ul style="list-style-type: none"> • 0 = maps to discard action • 1 = maps to accept action
Indirection	1 byte	<ul style="list-style-type: none"> • 0 = egress • 1 = ingress 	<ul style="list-style-type: none"> • 0 = adds egress terms to the output filter • 1 = adds ingress terms to the input filter
Spare	1 byte	–	–

Table 30: 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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 0 = no mapping performed • From destination-address address entry added to term
Source IP prefix	1 byte	<ul style="list-style-type: none"> • 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 	<ul style="list-style-type: none"> • 0 = no mapping performed • From source-address prefix entry added to term
Destination IP prefix	1 byte	<ul style="list-style-type: none"> • 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 	<ul style="list-style-type: none"> • 0 = no mapping performed • From destination-address prefix entry added to term
Protocol	1 byte	Protocol type	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 0 = no compare • 1 = less than • 2 = equal to • 3 = greater than • 4 = not equal to 	<ul style="list-style-type: none"> • 0 = no mapping performed • 1 – 3 = mapped to corresponding option • 4 = mapped to except match option

Table 30: Ascend-Data-Filter Attribute Fields (*continued*)

Action or Classifier	Format	Value	Junos OS Filter Function
Destination port qualifier	1 byte	<ul style="list-style-type: none"> • 0 = no compare • 1 = less than • 2 = equal to • 3 = greater than • 4 = not equal to 	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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.

Table 30: Ascend-Data-Filter Attribute Fields (*continued*)

Action or Classifier	Format	Value	Junos OS Filter Function
Rate-limit profile	1–41 bytes	<ul style="list-style-type: none"> • 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 <code>[edit firewall]</code> stanza. 	Maps to the policer <i>policer-name</i> action modifier of the same name. The action policer <i>name</i> is added to term.

RELATED DOCUMENTATION

[Ascend-Data-Filter Policies for Subscriber Management Overview](#) | 303

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
```

RELATED DOCUMENTATION

[Ascend-Data-Filter Policies for Subscriber Management Overview | 303](#)

[Ascend-Data-Filter Attribute Fields | 305](#)

[Verifying and Managing Dynamic Ascend-Data-Filter Policy Configuration | 320](#)

[Example: Configuring Dynamic Ascend-Data-Filter Support for Subscriber Access | 310](#)

[Example: Configuring Static Ascend-Data-Filter Support for Subscriber Access | 315](#)

Example: Configuring Dynamic Ascend-Data-Filter Support for Subscriber Access

IN THIS SECTION

- [Requirements | 311](#)
- [Overview | 311](#)
- [Configuration | 311](#)
- [Verification | 313](#)

This example shows how to configure support for dynamic Ascend-Data-Filter policies.

Requirements

- Ensure that the Ascend-Data-Filter has been configured on the RADIUS server.
- Create the dynamic profile. See *Dynamic Profiles Overview*.
- Configure RADIUS support. See *RADIUS Servers and Parameters for Subscriber Access*.

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
```

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.

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 | 303](#)

[Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions | 308](#)

Example: Configuring Static Ascend-Data-Filter Support for Subscriber Access

IN THIS SECTION

- [Requirements | 315](#)
- [Overview | 316](#)
- [Configuration | 316](#)
- [Verification | 318](#)

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

- Create the dynamic profile. See *Dynamic Profiles Overview*.
- Configure RADIUS support. See *RADIUS Servers and Parameters for Subscriber Access*.

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 CB007100 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
...
adf-profile-v4 {
  interfaces {
    "$junos-interface-ifd-name" {
      unit "$junos-underlying-interface-unit" {
        family inet {
          filter {
            adf {
              rule "01000100 CB007100 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 203.0.113.0 with wildcard mask 255.255.255.0 to any destination.

Table 31 on page 317 lists the values specified in the Ascend-Data-Filter rule.

Table 31: Ascend-Data-Filter Rule

Action or Classifier	Hex Value	Junos OS Filter Function
Type	01	IPv4

Table 31: Ascend-Data-Filter Rule (continued)

Action or Classifier	Hex Value	Junos OS Filter Function
Forward	00	Forward
Indirection	01	Ingress
Spare	00	None
Source IP address	CB007100	203.0.113.0
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
Destination port qualifier	00	None
Reserved	0000	None

Verification

IN THIS SECTION

- [Verifying that Static Ascend-Data-Filter Rules are Applied to Subscriber Sessions | 319](#)
- [Verifying Static Ascend-Data-Filter Usage | 319](#)

To confirm that the configuration is working properly, perform these tasks:

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: 00:00:5E:00:53:01
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: 01000100CB00710000000000018000000000000000000000
                                from {
                                    destination-address 203.0.113.0/24;
                                }
                                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.
- 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 | 303](#)

[Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions | 308](#)

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](#) | 303

[Dynamically Applying Ascend-Data-Filter Policies to Subscriber Sessions](#) | 308

Configuring Fast Update Filters to Provide More Efficient Processing Over Classic Static Filters

IN THIS CHAPTER

- [Fast Update Filters Overview | 323](#)
- [Basic Fast Update Filter Syntax | 327](#)
- [Configuring Fast Update Filters | 328](#)
- [Example: Configuring Fast Update Filters for Subscriber Access | 330](#)
- [Match Conditions and Actions in Fast Update Filters | 331](#)
- [Configuring the Match Order for Fast Update Filters | 333](#)
- [Fast Update Filter Match Conditions | 334](#)
- [Fast Update Filter Actions and Action Modifiers | 335](#)
- [Configuring Terms for Fast Update Filters | 336](#)
- [Configuring Filters to Permit Expected Traffic | 337](#)
- [Avoiding Conflicts When Terms Match | 338](#)
- [Associating Fast Update Filters with Interfaces in a Dynamic Profile | 345](#)

Fast Update Filters Overview

IN THIS SECTION

- [Fast Update Filter Components | 324](#)
- [Fast Update Filter Processing | 325](#)
- [Fast Update Filter Names | 326](#)
- [Guidelines for Creating and Applying Fast Update Filters | 326](#)

Fast update filters provide more efficient filter processing over classic static filters when dynamic services are implemented for multiple subscribers that share the same logical interface.

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 the router to 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 328](#).
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 345](#).
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 331](#).

This overview covers:

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 334](#) 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 335](#) 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 | 335](#)

[Fast Update Filter Match Conditions | 334](#)

[Avoiding Conflicts When Terms Match | 338](#)

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 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:

- Dynamic application of input and output filters is supported.
- 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 337](#).
- 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 333](#).
- 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

[Understanding Dynamic Firewall Filters | 226](#)

[Classic Filters Overview | 229](#)

[Dynamically Attaching Statically Created Filters for Any Interface Type | 246](#)

[Dynamically Attaching Statically Created Filters for a Specific Interface Family Type | 245](#)

[Verifying and Managing Firewall Filter Configuration | 389](#)

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](#) | 328

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.

```
[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 333](#).

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 336](#).

RELATED DOCUMENTATION

[Configuring the Match Order for Fast Update Filters](#) | 333

[Configuring Terms for Fast Update Filters | 336](#)

[Associating Fast Update Filters with Interfaces in a Dynamic Profile | 345](#)

[Fast Update Filters Overview | 323](#)

[Dynamic Profiles Overview](#)

[Guidelines for Configuring Firewall Filters](#)

[Guidelines for Applying Standard Firewall Filters](#)

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 | 328](#)

Match Conditions and Actions in Fast Update Filters

IN THIS SECTION

- [Match Conditions | 331](#)
- [Actions | 332](#)
- [Adding Terms Only Once | 332](#)

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

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 334](#) 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.

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 335](#) 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 | 336](#)

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]  
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 Terms for Fast Update Filters 336
Fast Update Filters Overview 323
Dynamic Profiles Overview
Fast Update Filter Match Conditions 334
Guidelines for Configuring Firewall Filters

Fast Update Filter Match Conditions

Table 32: Fast Update Filter Match Conditions

Match Condition	Description
destination-address <i>prefix</i>	IP destination address field.
destination-port <i>number</i>	TCP or UDP destination port field. Can be a single number, a single range, or one of the standard port synonyms.
dscp <i>number</i>	Differentiated services code point. Can be a single number, a single range, or the standard synonyms. IPv4 only.
match-terms <i>string-of-conditions</i>	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 <i>number</i>	IP protocol field. Can be a single number, a single range, or one of the standard protocol synonyms. IPv4 only.
source-address <i>prefix</i>	IP source address field.
source-port <i>number</i>	TCP or UDP source port field. Can be a single number, a single range, or one of the standard protocol synonyms.

RELATED DOCUMENTATION

Configuring Fast Update Filters 328

Fast Update Filter Actions and Action Modifiers

Table 33: Fast Update Filter Actions and Action Modifiers

Action or Action Modifier	Description
Actions	
accept	Accept the packet.
action-terms <i>string-of-actions</i>	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 <i>routing-instance</i>	Forward packets to specified routing instance.
Action Modifiers	
count <i>counter-name</i>	Increment the specified counter.
forwarding-class <i>class</i>	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 <i>policer-name</i>	Rate-limit packets based on the specified policer.

RELATED DOCUMENTATION

[Configuring Fast Update Filters](#) | 328

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 334](#) 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 335](#) 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 335](#) 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 | 328](#)

[Configuring the Match Order for Fast Update Filters | 333](#)

[Fast Update Filters Overview | 323](#)

[Fast Update Filter Match Conditions | 334](#)

[Fast Update Filter Actions and Action Modifiers | 335](#)

[Stateless Firewall Filter Overview](#)

[Stateless Firewall Filter Components](#)

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 333](#). 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 | 333](#)

[Configuring Terms for Fast Update Filters | 336](#)

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 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	203.0.113.2/32	http	count t55_cntr accept
t999	subscriber's address	203.0.113.2/32	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 203.0.113.2/32;
          destination-port http;
        }
        then {
          count t55_cntr;
          accept;
        }
      }
      term t999 {
        from {
          source-address $junos-subscriber-ip-address;

```

```
        destination-address 203.0.113.2/32;
        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	203.0.113.2/32	any (wildcard)	count t3_cntr accept
t55	subscriber's address	203.0.113.2/32	http	count t55_cntr accept

Term	source-address	destination-address	destination-port	Action
t999	subscriber's address	203.0.113.2/32	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 203.0.113.2/32;
        }
        then {
          count t3_cntr;
          accept;
        }
      }
      term t55 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 203.0.113.2/32;
          destination-port http;
        }
        then {
          count t55_cntr;
          accept;
        }
      }
      term t999 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 203.0.113.2/32;

```

```
        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	203.0.113.2/32	0-80	count c1_cntr accept
src1	subscriber's address	203.0.113.2/32	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 203.0.113.2/32;
          destination-port 0-80;
        }
        then {
          count c1_cntr;
          accept;
        }
      }
      term src1 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 203.0.113.2/32;
          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	203.0.113.2/32	udp	any (wildcard)	count c20_cntr accept
src21	subscriber's address	203.0.113.2/32	any (wildcard)	http	count c21_cntr accept

Term	source-address	destination-address	protocol	destination-port	Action
src21	subscriber's address	203.0.113.2/32	any (wildcard)	https	count c22_cntr accept

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 203.0.113.2/32;
          protocol udp;
        }
        then {
          count c20_cntr;
          accept;
        }
      }
      term src21 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 203.0.113.2/32;
          destination-port http;
        }
        then {
          count c21_cntr;
          accept;
        }
      }
      term src22 {
        from {
          source-address $junos-subscriber-ip-address;
          destination-address 203.0.113.2/32;
          destination-port https;
        }
        then {

```



```

        count c22_cntr;
        accept;
    }
}

```

RELATED DOCUMENTATION

[Configuring Fast Update Filters | 328](#)

[Configuring Terms for Fast Update Filters | 336](#)

[Configuring the Match Order for Fast Update Filters | 333](#)

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

[Fast Update Filters Overview](#) | 323

Guidelines for Configuring Firewall Filters

Guidelines for Applying Standard Firewall Filters

Defending Against DoS and DDoS Attacks Using Unicast RPF and Fail Filters

IN THIS CHAPTER

- [Unicast RPF in Dynamic Profiles for Subscriber Interfaces | 347](#)

Unicast RPF in Dynamic Profiles for Subscriber Interfaces

IN THIS SECTION

- [Unicast RPF in Dynamic Profiles for Subscriber Interfaces | 347](#)
- [Configuring Unicast RPF in Dynamic Profiles for Subscriber Interfaces | 348](#)
- [Configuring a Fail Filter for Unicast RPF in Dynamic Profiles for Subscriber Interfaces | 349](#)
- [Example: Configuring Unicast RPF in a Dynamic Profile on MX Series Routers | 350](#)

Unicast RPF in Dynamic Profiles for Subscriber Interfaces

Unicast reverse-path forwarding (RPF) provides a way to reduce the effect of denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on IPv4 and IPv6 interfaces. When you configure unicast RPF on an interface, it checks the packet source address. Packets that pass the check are forwarded. Packets that fail the check are dropped, or if a fail filter is configured, are passed to the filter for further evaluation.

Unicast RPF has two behavioral modes, *strict* and *loose*. When you configure unicast RPF in a dynamic profile, strict mode is the default. In strict mode, unicast RPF checks whether the source address of the incoming packet matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix. In loose mode, unicast RPF checks only whether the source address has a match in the routing table. It does not check whether the interface expects to receive a packet from a specific source address.

For both modes, when an incoming packet fails the unicast RPF check, the packet is not accepted on the interface. Instead, unicast RPF counts the packet and sends it to an optional fail filter, if present. The fail filter determines what further action is taken on the packet. In the absence of a fail filter, the packet is silently discarded.

Starting in Junos OS Release 19.1R1, the **show interfaces statistics *logical-interface-name* detail** command displays unicast RPF statistics for dynamic logical interfaces when either **rpf-check** or **rpf-check mode loose** is enabled on the interface. No additional statistics are displayed when **rpf-check fail-filter *filter-name*** is configured on the interface. The **clear interfaces statistics *logical-interface-name*** command clears RPF statistics.

Configuring Unicast RPF in Dynamic Profiles for Subscriber Interfaces

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 a unicast RPF check 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. Enable the RPF check in strict or loose mode.

- Configure strict mode to check whether the source address of the incoming packet matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix:

```
[edit dynamic-profiles profile-name interface interface-name unit logical-unit-number family inet]
user@host# set rpf-check
```

- Configure loose mode to check only whether the source address has a match in the routing table:

```
[edit dynamic-profiles profile-name interface interface-name unit logical-unit-number family inet]
user@host# set rpf-check mode loose
```

4. (Optional except for DHCP) Enable 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
```

For information about defining a fail filter, see [“Configuring a Fail Filter for Unicast RPF in Dynamic Profiles for Subscriber Interfaces”](#) on page 349.

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.

```
[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.

Example: Configuring Unicast RPF in a Dynamic Profile on MX Series Routers

IN THIS SECTION

- [Requirements | 350](#)
- [Overview | 351](#)
- [Configuration | 352](#)
- [Verification | 356](#)

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

This example uses the following software and hardware components:

- An MX Series 5G Universal Routing Platform

Before you begin:

- Configure the dynamic profile that you intend to use to apply the RPF check.

See *Configuring a Basic Dynamic Profile*.

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 *Understanding Unicast RPF (Routers)* 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.

Configuration

IN THIS SECTION

- [Configuring the Dynamic Profile to Apply RPF Checking to Dynamic VLAN Demux Interfaces | 352](#)
- [Configuring the RPF-Check Fail Filter | 353](#)

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

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

```

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

IN THIS SECTION

- [Verifying That Unicast RPF Is Enabled on the Router | 356](#)

To confirm that the configuration is correct, perform these tasks:

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.

Release History Table

Release	Description
19.1R1	Starting in Junos OS Release 19.1R1, the show interfaces statistics <i>logical-interface-name</i> detail command displays unicast RPF statistics for dynamic logical interfaces when either rpf-check or rpf-check mode loose is enabled on the interface.

RELATED DOCUMENTATION

For more detailed information about unicast RPF in general, see *Understanding Unicast RPF (Routers)*

Improving Scaling and Performance of Filters on Static Subscriber Interfaces

IN THIS CHAPTER

- [Firewall Filters and Enhanced Network Services Mode Overview | 359](#)
- [Configuring a Filter for Use with Enhanced Network Services Mode | 362](#)

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.

NOTE: Do not use enhanced mode for firewall filters that are intended for control plane traffic. Control plane filtering is handled by the Routing Engine kernel, which cannot use the term-based format of the enhanced mode filters.

Table 34 on page 360 shows the configuration options when determining enhanced network services mode usage.

Table 34: 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. Only term-based (enhanced) firewall filters will be generated, regardless of the setting of the **enhanced-mode** statement at the **[edit chassis network-services]** hierarchy level, if any of the following are true:

- Flexible filter match conditions are configured at the **[edit firewall family *family-name* filter *filter-name* term *term-name* from]** or **[edit firewall filter *filter-name* term *term-name* from]** hierarchy levels.
- A tunnel header push or pop action, such as GRE encapsulate or decapsulate is configured at the **[edit firewall family *family-name* filter *filter-name* term *term-name* then]** hierarchy level.
- Payload-protocol match conditions are configured at the **[edit firewall family *family-name* filter *filter-name* term *term-name* from]** or **[edit firewall filter *filter-name* term *term-name* from]** hierarchy levels.
- An extension-header match is configured at the **[edit firewall family *family-name* filter *filter-name* term *term-name* from]** or **[edit firewall filter *filter-name* term *term-name* from]** hierarchy levels.
- A match condition is configured that only works with MPC cards, such as firewall bridge filters for IPv6 traffic.



WARNING: Any firewall filter meeting the previous criteria will not be applied to the loopback, lo0, interface of DPC based FPCs. This means that term-based (enhanced) filters configured for use on the loopback interface of a DPC based FPC will not be applied. This will leave the RE unprotected by that filter.

RELATED DOCUMENTATION

[Network Services Mode Overview](#)

[Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers](#)

[Configuring a Filter for Use with Enhanced Network Services Mode](#) | 362

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 359](#) for details.

To configure a stateless firewall filter to use enhanced mode:

1. Create or edit the stateless firewall filter.

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 [Example: Configuring and Applying a Simple Filter](#) for a filter configuration example.

RELATED DOCUMENTATION

Understanding How to Use Standard Firewall Filters

Network Services Mode Overview

[Firewall Filters and Enhanced Network Services Mode Overview | 359](#)

Configuring Junos OS to Run a Specific Network Services Mode in MX Series Routers

[Understanding Dynamic Firewall Filters | 226](#)

Configuring Dynamic Service Sets

IN THIS CHAPTER

- [Dynamic Service Sets Overview | 365](#)
- [Associating Service Sets with Interfaces in a Dynamic Profile | 366](#)
- [Verifying and Managing Service Sets Information | 367](#)

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 | 366](#)

[Verifying and Managing Service Sets Information | 367](#)

[Understanding Service Sets](#)

[Applying Filters and Services to Interfaces](#)

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
```

RELATED DOCUMENTATION

[Dynamic Service Sets Overview | 365](#)

[Verifying and Managing Service Sets Information | 367](#)

[Configuring Service Sets to be Applied to Services Interfaces](#)

[Applying Filters and Services to Interfaces](#)

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 | 365](#)

[Associating Service Sets with Interfaces in a Dynamic Profile | 366](#)

[CLI Explorer](#)

Configuring Rate-Limiting Premium and Non-Premium Traffic on an Interface Using Hierarchical Policers

IN THIS CHAPTER

- [Methods for Regulating Traffic by Applying Hierarchical Policers | 369](#)
- [Hierarchical Policar Applied as Filter Action | 372](#)
- [Example: Configuring Hierarchical Policers to Limit Rates of Services in a Static Environment | 373](#)

Methods for Regulating Traffic by Applying Hierarchical Policers


You can deploy policers to enforce service level agreements limiting the input rate at the edge, and at the boundary between domains, to guarantee an equitable deployment of the service among the different domains. Policers determine whether each packet conforms (falls within the traffic contract), exceeds (using up the excess burst capacity), or violates (totally out of the traffic contract rate) the configured traffic policies, and then sets the prescribed action.

Hierarchical policers rate-limit premium traffic separately from the aggregate traffic on an interface as determined by different configured rates. You can use a hierarchical policer to rate-limit ingress Layer 2 traffic at a physical or logical interface and apply different policing actions based on whether the traffic or packets are classified for expedited forwarding (EF) or for a lower priority, such as non-expedited forwarding (non-EF).

Hierarchical policers provide cross-functionality between the configured physical interface and the Packet Forwarding Engine. You can apply a hierarchical policer for premium and aggregate (premium plus normal) traffic levels to a logical interface.

Hierarchical policing uses two token buckets, one for premium (EF) traffic and one for aggregate (non-EF) traffic, as shown in [Figure 6 on page 370](#).

Figure 6: Hierarchical Policer



The class-of-service (CoS) configuration determines which traffic is EF and which is non-EF. Logically, hierarchical policing is achieved by chaining two policers.

- **Premium policer**—You configure the premium policer with traffic limits for high-priority EF traffic only: a guaranteed bandwidth and a corresponding burst-size limit. EF traffic is categorized as nonconforming when its average arrival rate exceeds the guaranteed bandwidth and its average packet size exceeds the premium burst-size limit. For a premium policer, the only configurable action for nonconforming traffic is to discard the packets.
- **Aggregate policer**—You configure the aggregate policer (also known as a logical interface policer) with an aggregate bandwidth (to accommodate both high-priority EF traffic up to the guaranteed bandwidth and normal-priority non-EF traffic) and a burst-size limit for non-EF traffic only. Non-EF traffic is categorized as nonconforming when its average arrival rate exceeds the amount of aggregate bandwidth not currently consumed by EF traffic and its average packet size exceeds the burst-size limit defined in the aggregate policer. For an aggregate policer, the configurable actions for nonconforming traffic are to discard the packets, assign a forwarding class, or assign a packet loss priority (PLP) level.

NOTE: You must configure the bandwidth limit of the premium policer at or below the bandwidth limit of the aggregate policer. If the two bandwidth limits are equal, then only non-EF traffic passes through the interface unrestricted; no EF traffic arrives at the interface.

Ingress traffic is first classified into EF and non-EF traffic prior to applying a policer. EF traffic is guaranteed the bandwidth specified as the premium bandwidth limit, while non-EF traffic is rate-limited to the amount of aggregate bandwidth not currently consumed by the EF traffic. Non-EF traffic is rate-limited to the entire aggregate bandwidth only while no EF traffic is present.

Hierarchical policing uses two token buckets, one for aggregate (non-EF) traffic and one for premium (EF) traffic. In [Figure 6 on page 370](#), the premium policer policies EF traffic and the aggregate policer polices non-EF traffic. In the sample configuration that follows, the hierarchical policer is configured with the following components:

- Premium policer has a bandwidth limit set to 2 Mbps, burst-size limit set to 50 KB, and nonconforming action set to discard packets.
- Aggregate policer has a bandwidth limit set to 10 Mbps, burst-size limit set to 100 KB, and nonconforming action set to mark high PLP.

```
[edit]
user@host# show dynamic-profiles firewall
hierarchical-policer policer-agg-prem {
  aggregate {
    if-exceeding {
      bandwidth-limit 10m;
      burst-size-limit 100k;
    }
    then {
      loss-priority high;
    }
  }
  premium {
    if-exceeding {
      bandwidth-limit 2m;
      burst-size-limit 50k;
    }
    then {
      discard;
    }
  }
}
```

EF traffic is guaranteed a bandwidth of 2 Mbps. Bursts of EF traffic—EF traffic that arrives at the interface at rates above 2 Mbps—can also pass through the interface, provided that sufficient tokens are available in the 50 KB burst bucket. When no tokens are available, EF traffic is rate-limited using the discarded action associated with the premium policer.

Non-EF traffic is metered to a bandwidth limit that ranges between 8 Mbps and 10 Mbps, depending on the average arrival rate of the EF traffic. Bursts of non-EF traffic—non-EF traffic that arrives at the interface at rates above the current limit for non-EF traffic—also pass through the interface, provided that sufficient tokens are available in the 100 KB bandwidth bucket. Aggregate traffic in excess of the currently configured bandwidth or burst size are rate-limited using the action specified for the aggregate policer, which in this example is set to a high PLP.

The premium traffic is policed by both the premium policer and aggregate policer. Although the premium policer rate-limits the premium traffic, the aggregate policer decrements the credits but does not drop the packets. The aggregate policer rate-limits the non-premium traffic. 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

[Example: Configuring Hierarchical Policers to Limit Rates of Services in a Static Environment | 373](#)
[Hierarchical Policar Applied as Filter Action | 372](#)

Hierarchical Policar Applied as Filter Action

After you define firewall filters and policers, you must apply them to take effect.

- You can apply the same firewall filter to multiple interfaces at the same time. By default on MX Series routers, these filters aggregate their counters and policing actions when those interfaces share a Packet Forwarding Engine. To override this behavior and make each counter or policer function specific to each interface application, include the **interface-specific** statement in the firewall filter.

```
[edit dynamic-profiles profile-name firewall family family filter filter-name
user@host# set interface-specific
```

Interface-specific filters are particularly useful for IPTV services where television services are delivered using the IP suite over a packet-switched network instead of being delivered through traditional satellite signal and cable television formats.

NOTE: When you define an interface-specific filter, you must limit the filter name to no more than 52 bytes. Firewall filter names are restricted to 64 bytes in length and interface-specific filters have the specific-name appended to them to differentiate their counters and policing actions. If the automatically generated filter instance name exceeds this maximum length, the system may reject the filter's instance name.

- Alternatively, you can apply a policer to a logical interface either directly or indirectly through a filter that references the policer function. By default, policers are *term-specific*. Junos OS creates a separate policer instance when the same policer is referenced in multiple terms of a firewall filter.

Hierarchical policers provide cross-functionality between the configured physical interface and the Packet Forwarding Engine for provider edge applications. You can apply a hierarchical policer as a filter action for

premium and aggregate (premium plus normal) traffic levels to a logical interface. Additionally, an interface-specific filter can have a hierarchical policer as a filter action whether or not the hierarchical policer is a logical interface policer.

A logical interface policer (also known as an aggregate policer) can police the traffic from multiple protocol families without requiring a separate instantiation of a policer for each such family on the logical interface. You define a logical interface policer by including the **logical-interface-policer** statement when defining the policer.

```
[edit dynamic-profiles profile-name firewall policer policer-name
user@host# set logical-interface-policer
```

To apply a logical interface policer on an MX Series router as an action in a firewall filter term, you must specify both the **interface-specific** statement in the firewall filter and the **logical-interface-policer** statement in the related policer. Using a filter to evoke a logical interface filter has the added benefits of increased match flexibility as well as support for two-color policer styles (a policer that classifies traffic into two groups using only the **bandwidth-limit** and **burst-size-limit** parameters), which can only be attached at the family level through a filter action.

NOTE: A non-interface-specific filter can only have a hierarchical policer if no logical interface-specific filter action is specified.

RELATED DOCUMENTATION

[Methods for Regulating Traffic by Applying Hierarchical Policers | 369](#)

[Example: Configuring Hierarchical Policers to Limit Rates of Services in a Static Environment | 373](#)

Example: Configuring Hierarchical Policers to Limit Rates of Services in a Static Environment

IN THIS SECTION

- [Requirements | 374](#)
- [Overview | 374](#)

- Configuration | 375
- Verification | 385

This example shows how to configure a hierarchical policer and apply the policer to ingress Layer 2 traffic at a logical interface on an MX Series router.

Requirements

Before you begin, be sure that your environment meets the following requirements:

- The interface on which you apply the hierarchical policer is an interface hosted on an MX Series router.
- No other policer is applied to the input of the interface on which you apply the hierarchical policer.
- You are aware that, if you apply the hierarchical policer to logical interface on which an input filter is also applied, the policer is executed first.

Overview

In this example, you configure a hierarchical policer and apply the policer to ingress Layer 2 traffic at a logical interface. [Table 35 on page 374](#) describes the hierarchy levels at which you can configure and apply hierarchical policers on logical and physical interfaces.

Table 35: Hierarchical Policer Configuration and Application Summary

Policer Configuration	Layer 2 Application	Key Points
-----------------------	---------------------	------------

Hierarchical Policer

Hierarchically rate-limits Layer 2 ingress traffic for all protocol families. Cannot be applied to egress traffic, Layer 3 traffic, or at a specific protocol level of the interface hierarchy. Supported on interfaces on Dense Port Concentrators (DPCs) in MX Series routers.

<p>Aggregate and premium policing components of a hierarchical policer:</p> <pre>[edit dynamic-profiles profile-name firewall] hierarchical-policer policer-name { aggregate {</pre>	<p>Option A (physical interface)—Apply directly to Layer 2 input traffic on a physical interface:</p> <pre>[edit dynamic-profiles profile-name interfaces] interface-name { layer2-policer { input-hierarchical-policer policer-name; }</pre>	<p>Hierarchically rate-limit Layer 2 ingress traffic for all protocol families and logical interfaces configured on a physical interface.</p> <p>Include the layer2-policer configuration statement at the [edit dynamic-profiles</p>
--	---	---

Table 35: Hierarchical Policer Configuration and Application Summary (*continued*)

Policer Configuration	Layer 2 Application	Key Points
<pre> if-exceeding { bandwidth-limit <i>bps</i>; burst-size-limit <i>bytes</i>; } then { discard; forwarding-class <i>class-name</i>; loss-priority <i>supported-value</i>; } } premium { if-exceeding { bandwidth-limit <i>bps</i>; burst-size-limit <i>bytes</i>; } then { discard; } } } </pre>	<pre> } </pre>	<p><i>profile-name</i> interfaces <i>interface-name</i> hierarchy level.</p> <p>NOTE: If you apply a hierarchical policer at a physical interface, you cannot also apply a hierarchical policer to any of the member logical interfaces.</p>
	<p>Option B (logical interface)—Apply directly to Layer 2 input traffic on a logical interface:</p> <pre> [edit dynamic-profiles <i>profile-name</i> interfaces] <i>interface-name</i> { unit <i>unit-number</i> { layer2-policer { input-hierarchical-policer <i>policer-name</i>; } } } </pre>	<p>Hierarchically rate-limit Layer 2 ingress traffic for all protocol families configured on a specific logical interface.</p> <p>Include the layer2-policer configuration statement at the [edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>unit-number</i>] hierarchy level.</p> <p>NOTE: You must configure at least one protocol family for the logical interface.</p>

You apply the policer to the Gigabit Ethernet logical interface ge-1/2/0.0, which you configure for IPv4 traffic. When you apply the hierarchical policer to the logical interface, IPv4 traffic is hierarchically rate-limited. If you choose to apply the hierarchical policer to physical interface ge-1/2/0, hierarchical policing applies to IPv4 traffic across the logical interface as well.

Configuration

IN THIS SECTION

- [Configuring a Basic Dynamic Profile for Subscriber Management | 377](#)
- [Configuring the Interfaces | 378](#)
- [Configuring the Firewall Filter | 379](#)
- [Configuring the Forwarding Classes | 381](#)

- [Configuring the Hierarchical Policier | 382](#)
- [Applying the Hierarchical Policier to Layer 2 Ingress Traffic at a Physical or Logical Interface | 383](#)

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*.

To configure this example, perform the following tasks:

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set dynamic-profiles basic-profile
set dynamic-profiles basic-profile interfaces "$junos-interface-ifd-name"
set dynamic-profiles basic-profile interfaces "$junos-interface-ifd-name" unit "$junos-underlying-interface-unit"
set dynamic-profiles basic-profile interfaces "$junos-interface-ifd-name" unit $junos-underlying-interface-unit
  family inet
set dynamic-profiles interfaces ge-1/2/0 unit 0 family inet address 203.0.113.80/31
set dynamic-profiles basic-profile firewall family inet filter hierarch-filter
set dynamic-profiles basic-profile firewall family inet filter hierarch-filter interface-specific
set dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip1
set dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip2
set dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip1 from precedence
  critical-ecp protocol
set dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip1 from protocol tcp
set dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip1 then hierarchical-policer
  hp1-share filter-specific
set dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip2 from precedence
  internet-control
set dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip2 from protocol tcp
set dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip2 then hierarchical-policer
  hp2-share
set class-of-service forwarding-classes class fc0 queue-num 0 priority high policing-priority premium
set class-of-service forwarding-classes class fc1 queue-num 1 priority low policing-priority normal
set class-of-service forwarding-classes class fc2 queue-num 2 priority low policing-priority normal
set class-of-service forwarding-classes class fc3 queue-num 3 priority low policing-priority normal
set dynamic-profiles basic-profile firewall hierarchical-policer policer-agg-prem aggregate if-exceeding
  bandwidth-limit 10m burst-size-limit 100k
set dynamic-profiles basic-profile firewall hierarchical-policer policer-agg-prem aggregate then forwarding-class
  fc1
```

```
set dynamic-profiles basic-profile firewall hierarchical-policer policer-agg-prem premium if-exceeding
bandwidth-limit 2m burst-size-limit 50k
set dynamic-profiles basic-profile firewall hierarchical-policer policer-agg-prem premium then discard
set dynamic-profiles basic-profile interfaces ge-1/2/0 unit 0 layer2-policer input-hierarchical-policer
policer-agg-prem
```

Configuring a Basic Dynamic Profile for Subscriber Management

Step-by-Step Procedure

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. 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.

1. Create the new dynamic profile.

```
[edit]
user@host# set dynamic-profiles basic-profile
```

2. Define the ***interface-name*** variable 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# set interfaces "$junos-interface-ifd-name"
```

3. Define the ***variable-interface-name* 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. Define the family address type (inet for IPv4) for the **\$junos-interface-unit** variable.

```
[edit dynamic-profiles basic-profile interfaces "$junos-interface-ifd-name" unit $junos-underlying-interface-unit]
user@host# set family inet
```

Results

Confirm the configuration of the dynamic profile by entering the **show dynamic-profiles** 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 dynamic-profiles
dynamic-profiles {
  basic-profile {
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-underlying-interface-unit" {
          family inet;
        }
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Interfaces

Step-by-Step Procedure

Define the physical and logical interfaces for this hierarchical policer example.

1. Configure the physical interface.

```
[edit dynamic-profiles basic-profile]
user@host# set interfaces ge-1/2/0
```

2. Configure the logical interface as unit 0 with its IPv4 (inet) protocol family interface.

```
[edit dynamic-profiles basic-profile interfaces ge-1/2/0]
user@host# set unit 0 family inet address 203.0.113.80/31
```

NOTE: If you apply a Layer 2 policer to this logical interface, you must configure at least one protocol family.

Results

Confirm the configuration by entering the **show dynamic-profiles basic-profile 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 dynamic-profiles basic-profile interfaces
ge-1/2/0 {
  unit 0 {
    family inet {
      address 203.0.113.80/31;
    }
  }
}
```

Configuring the Firewall Filter

Step-by-Step Procedure

To configure a hierarchical policer as a filter action, you must first configure a firewall filter.

1. Configure the family address type (inet for IPv4) for the firewall filter and specify the filter name.

We recommend that you name the filter something that indicates the filter's purpose.

```
[edit dynamic-profiles basic-profile]
user@host# set firewall family inet filter hierarch-filter
```

2. To override the aggregation of the counters and policing actions and make each counter or policy function specific to each interface application, include the **interface-specific** statement in the filter.

```
[edit dynamic-profiles basic-profile firewall family inet filter hierarch-filter]
user@host# set interface-specific
```

3. Specify the term names for the filter.

Make each term name unique and represent what its function is.


```
[edit dynamic-profiles basic-profile firewall family inet filter hierarch-filter]
user@host# set term match-ip1
user@host# set term match-ip2
```

4. In each firewall filter term, specify the conditions used to match components of a packet.

Configure the first term to match IPv4 packets received through TCP with the IP precedence field critical-ecp (0xa0) protocol, and apply the hierarchical policer as a filter action.

```
[edit dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip1]
user@host# set from precedence critical-ecp protocol
user@host# set from protocol tcp
```

5. Specify the actions to take when the packet matches all of the conditions in the first term. Enable all hierarchical policers in one filter to share the same policer instance in the Packet Forward Engine.

```
[edit dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip1]
user@host# set then hierarchical-policer hp1-share filter-specific
```

6. Configure the second term to match IPv4 packets received through TCP with the IP precedence field internet-control (0xc0), and apply the hierarchical policer as a filter action.

```
[edit dynamic-profiles basic-profile firewall family inet filter hierarch-filter term match-ip2]
user@host# set from precedence internet-control
user@host# set from protocol tcp
```

7. Specify the actions to take when the packet matches all of the conditions in the second term.

```
[edit dynamic-profiles basic-profile firewall family inet filter inet-filter term match-ip2]
user@host# set then hierarchical-policer hp2-share
```

Results

Confirm the configuration by entering the **show dynamic-profiles basic-profile 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 dynamic-profiles basic-profile firewall
```

```

family inet {
  filter hierarch-filter {
    interface-specific;
    term match-ip1 {
      from {
        precedence critical-ecp protocol;
        protocol tcp;
      }
      then hierarchical-policer hp1-share;
    }
    term match-ip2 {
      from {
        precedence internet-control;
        protocol tcp;
      }
      then hierarchical-policer hp2-share;
    }
  }
}

```

Configuring the Forwarding Classes

Step-by-Step Procedure

Define forwarding classes referenced as aggregate policer actions. For hierarchical policers to work, ingress traffic must be correctly classified into premium and non-premium buckets. Some class-of-service (CoS) configuration is required because the hierarchical policer must be able to separate premium/expedited forwarding (EF) traffic from non-premium/non-EF traffic.

1. Enable configuration of the forwarding classes.

```

[edit]
user@host# set class-of-service forwarding-classes

```

2. Define CoS forwarding classes to include the designation of which forwarding class is premium. This defaults to the forwarding class associated with EF traffic.

```

[edit class-of-service forwarding-classes]
user@host# set class fc0 queue-num 0 priority high policing-priority premium
user@host# set class fc1 queue-num 1 priority low policing-priority normal
user@host# set class fc2 queue-num 2 priority low policing-priority normal
user@host# set class fc3 queue-num 3 priority low policing-priority normal

```

Results

Confirm the configuration of the forwarding classes referenced as aggregate policer actions by entering the **show class-of-service** 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 class-of-service
forwarding-classes {
  class fc0 queue-num 0 priority high policing-priority premium;
  class fc1 queue-num 1 priority low policing-priority normal;
  class fc2 queue-num 2 priority low policing-priority normal;
  class fc3 queue-num 3 priority low policing-priority normal;
}
```

Configuring the Hierarchical Policier

Step-by-Step Procedure

Configure the aggregate and premium policing components of a hierarchical policier.

1. Enable configuration of the hierarchical policier.

```
[edit dynamic-profiles basic-profile]
user@host# set firewall hierarchical-policier policer-agg-prem
```

2. Configure the aggregate policier to have a bandwidth limit set to 10 Mbps, burst-size limit set to 100 KB, and nonconforming action set to change the forwarding class to fc1.

```
[edit dynamic-profiles basic-profile firewall hierarchical-policier policer-agg-prem]
user@host# set aggregate if-exceeding bandwidth-limit 10m burst-size-limit 100k
user@host# set aggregate then forwarding-class fc1
```

NOTE: For aggregate policers, the configurable actions for a packet in a nonconforming flow are to discard the packet, change the loss priority, or change the forwarding class.

3. Configure the premium policier to have a bandwidth limit set to 2 Mbps, burst-size limit set to 50 KB, and nonconforming action set to discard packets.

```
[edit dynamic-profiles basic-profile firewall hierarchical-policier policer-agg-prem]
user@host# set premium if-exceeding bandwidth-limit 2m burst-size-limit 50k
user@host# set premium then discard
```

NOTE: The bandwidth limit for the premium policer must not be greater than that of the aggregate policer. For the premium policers, the only configurable action for a packet in a nonconforming traffic flow is to discard the packet.

Results

Confirm the configuration of the hierarchical policer by entering the **show dynamic-profiles basic-profile 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 dynamic-profiles basic-profile firewall
hierarchical-policer policer-agg-prem {
  aggregate {
    if-exceeding {
      bandwidth-limit 10m;
      burst-size-limit 100k;
    }
    then {
      forwarding-class fc1;
    }
  }
  premium {
    if-exceeding {
      bandwidth-limit 2m;
      burst-size-limit 50k;
    }
    then {
      discard;
    }
  }
}
```

Applying the Hierarchical Policer to Layer 2 Ingress Traffic at a Physical or Logical Interface

Step-by-Step Procedure

You can apply policers directly to an interface or applied through a filter to affect only matching traffic. In most cases, you can invoke a policing function at ingress, egress, or in both directions.

- For physical interfaces, a hierarchical policer uses a single policer instance to rate-limit all logical interfaces and protocol families configured on a physical interface, even if the logical interfaces have mutually exclusive families such as inet or bridge.
- For logical interfaces, a hierarchical policer can police the traffic from multiple protocol families without requiring a separate instantiation of a policer for each such family on the logical interface.

To hierarchically rate-limit Layer 2 ingress traffic for IPv4 traffic on logical interface ge-1/2/0.0, reference the policer from the logical interface configuration.

1. Configure the logical interface.

```
[edit dynamic-profiles basic-profile]
user@host# set interfaces ge-1/2/0 unit 0
```

When you apply a policer to Layer 2 traffic at a logical interface, you must define at least one protocol family for the logical interface.

2. Apply the policer to the logical interface.

```
[edit dynamic-profiles basic-profile interfaces ge-1/2/0 unit 0]
user@host# set layer2-policer input-hierarchical-policer policer-agg-prem
```

Alternatively, to hierarchically rate-limit Layer 2 ingress traffic for all protocol families and for *all logical interfaces* configured on physical interface ge-1/2/0, reference the policer from the physical interface configuration.

Results

Confirm the configuration of the hierarchical policer by entering the **show dynamic-profiles basic-profile 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 dynamic-profiles basic-profile interfaces
ge-1/2/0 {
  unit 0 {
    layer2-policer {
      input-hierarchical-policer policer-agg-prem;
    }
    family inet {
      address 203.0.113.80/31;
    }
  }
}
```

```

    }
  }
}

```

Verification

IN THIS SECTION

- [Displaying Traffic Statistics for the Interface | 385](#)
- [Displaying Number of Packets Policed by the Specified Policer | 388](#)

Confirm that the configuration is working properly.

Displaying Traffic Statistics for the Interface

Purpose

Verify the traffic flow through the physical interface.

Action

Use the **show interfaces** operational mode command for physical interface ge-1/2/0, and include the **detail** or **extensive** option.

```
user@host> show interfaces ge-1/2/0 extensive
```

```

Physical interface: ge-1/2/0, Enabled, Physical link is Down
  Interface index: 156, SNMP ifIndex: 630, Generation: 159
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 1000mbps, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags      : Present Running Down
  Interface flags: Hardware-Down 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: 00:00:5E:00:53:4c, Hardware address: 00:00:5E:00:53:4c

```

```

Last flapped   : 2014-11-10 13:36:25 EST (01:26:30 ago)
Statistics last cleared: Never

Traffic statistics:
  Input bytes   :                0                0 bps
  Output bytes  :               42                0 bps
  Input packets :                0                0 pps
  Output packets:                1                0 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: 0, 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, 8 in use

Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0                      0                0                0
  1                      0                0                0
  2                      0                0                0
  3                      0                0                0
  4                      0                0                0
  5                      0                0                0
  6                      0                0                0
  7                      0                0                0

Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
  4                be1
  5                ef1
  6                af1
  7                nc1

```

```

Active alarms : LINK
Active defects : LINK
MAC statistics:
    Receive          Transmit
Total octets          0          0
Total packets         0          0
Unicast packets       0          0
Broadcast packets     0          0
Multicast packets     0          0
CRC/Align errors      0          0
FIFO errors           0          0
MAC control frames    0          0
MAC pause frames      0          0
Oversized frames      0
Jabber frames         0
Fragment frames       0
VLAN tagged frames    0
Code violations        0
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   0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Incomplete
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       500000000     5          0      low
none
Interface transmit statistics: Disabled

```

Meaning

The command output section for **Traffic statistics** lists the number of bytes and packets received and transmitted on the interface.

Displaying Number of Packets Policed by the Specified Policer

Purpose

Verify the number of packets evaluated by the policer. Premium policer counters are not supported.

Action

Use the **show policer** operational mode command and optionally specify the name of the policer **policer-agg-prem**. The command output displays the number of packets evaluated by the specified policer in each direction.

user@host> **show policer policer-agg-prem**

Policers:		
Name	Bytes	Packets
policer-agg-prem-ge-1/2/0.0-inet-i	10372300	103723

The **-inet-i** suffix denotes a policer applied to IPv4 input traffic. In this example, the policer is applied to input traffic only.

Meaning

The command output displays the number of packets evaluated by the specified policer in each direction.

RELATED DOCUMENTATION

Methods for Regulating Traffic by Applying Hierarchical Policers 369
Hierarchical Policer Applied as Filter Action 372

Monitoring and Managing Firewalls for Subscriber Access

IN THIS CHAPTER

- [Verifying and Managing Firewall Filter Configuration | 389](#)
- [Enhanced Policer Statistics Overview | 390](#)

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 323](#) for information.

Action

- To display statistics for firewall filters:

```
user@host> show firewall
```

- To display firewall log information:

```
user@host> show firewall log
```

- To clear filter counters:

```
user@host> clear firewall all
```

RELATED DOCUMENTATION

[Classic Filters Overview | 229](#)

[Fast Update Filters Overview | 323](#)

[CLI Explorer](#)

Enhanced Policer Statistics Overview

You can use the enhanced policer statistics to analyze traffic for debugging purposes on MPC/MIC interfaces on MX Series routers and Multi-Rate Ethernet Enhanced Queuing IP Services DPC with SFP and XFP.

Enhanced policer statistics provide the following:

- Offered packet statistics for traffic subjected to policing.
- OOS packet statistics for packets that are marked out-of-specification by the policer. Changes to all packets that have out-of-specification actions, such as discard, color marking, or forwarding-class, are included in 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 within-specification statistics; when the policer action is non-discard (loss-priority or forwarding-class), the statistics are included in this counter.

RELATED DOCUMENTATION

[*show policer*](#)

[show firewall | 1378](#)

[enhanced-policer | 846](#)

4

PART

Configuring Dynamic Multicast

Configuring Dynamic IGMP to Support IP Multicasting for Subscribers | **393**

Configuring Dynamic MLD to Enable Subscribers to Access Multicast Networks | **401**

Configuring Dynamic IGMP to Support IP Multicasting for Subscribers

IN THIS CHAPTER

- [Dynamic IGMP Configuration Overview | 393](#)
- [Subscriber Management IGMP Model Overview | 394](#)
- [Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)
- [Example: IGMP Dynamic Profile | 397](#)
- [Configuring SSM Mapping for Dynamic IGMP and MLD | 399](#)

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.

Dynamic IGMP consists of a subset of the full range of IGMP capabilities available for static IGMP configuration, applied to dynamic interfaces by means of a dynamic profile. For detailed information about static IGMP configuration, see *Configuring IGMP*. Much of the static configuration documentation is directly applicable to dynamic IGMP. Note that the following statements that appear in the dynamic IGMP CLI hierarchy are configurable, but have no effect: **accounting**, **group-threshold**, **log-interval**, and **no-accounting**. These statements are not needed at a subscriber level, where typically no more than tens of joins are expected.

Refer to the *Multicast Protocols User Guide* for a comprehensive understanding of Junos OS support for multicast protocols.

RELATED DOCUMENTATION

[Dynamic Profiles Overview](#)[Subscriber Management IGMP Model Overview | 394](#)[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)[Configuring IGMP](#)

Subscriber Management IGMP Model Overview

In an IPTV network, channel changes occur when a set-top box (STB) sends IGMP commands that inform an upstream device (for example, a multiservice access node [MSAN] or services router) whether to start or stop sending multicast groups to the subscriber. In addition, IGMP hosts periodically request notification from the STB about which channels (multicast groups) are being received.

You can implement IGMP in the subscriber management network in the following ways:

- **Static IGMP**—All multicast channels are sent to the MSAN. When the MSAN receives an IGMP request to start or stop sending a channel, it adds the subscriber to the multicast group and then discards the IGMP packet.
- **IGMP Proxy**—Only multicast channels currently being viewed are sent to the MSAN. If the MSAN receives a request to view a channel that is not currently being forwarded to the MSAN, it forwards the request upstream. However, the upstream device does not see all channel change requests from each subscriber, limiting bandwidth control options.
- **IGMP Snooping**—Only multicast channels currently being viewed are sent to the MSAN. The MSAN forwards all IGMP requests upstream, unaltered, even if it is already receiving the channel. The upstream device sees all channel change requests from each subscriber. Using IGMP snooping enables the broadband services router to determine the mix of services and the bandwidth requirements of each subscriber and adjust the bandwidth made available to each service.
- **IGMP Passthrough**—The MSAN transparently passes IGMP packets upstream to the broadband services router.

IGMP hosts (sources) also periodically verify that they are sending the correct traffic by requesting that each client send information about what multicast groups it wants to receive. The responses to this *IGMP query* can result in a substantial upstream traffic burst.

IGMPv2 is the minimum level required to support IPTV, and is the most widely deployed. Emerging standards specify IGMPv3.

RELATED DOCUMENTATION

Configuring Dynamic DHCP Client Access to a Multicast Network

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*.

2. Configure the necessary router interfaces that you want accessing DHCP clients to use.

See *DHCP Subscriber Interface Overview* 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*.

4. Configure all RADIUS values that you want the profiles to use when validating DHCP clients for access to the multicast network.

See *RADIUS Servers and Parameters for Subscriber Access*.

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
Dynamic Profiles Overview

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 203.0.113.210;
  }
}
}
}
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 Dynamic DHCP Client Access to a Multicast Network](#) | 395

Configuring SSM Mapping for Dynamic IGMP and MLD

Source-specific multicast (SSM) is a service model that identifies session traffic by both source and group address. SSM builds shortest-path trees (SPTs) directly represented by (S,G) pairs. The “S” refers to the source’s unicast IP address, and the “G” refers to the specific multicast group address. The SSM (S,G) pairs are called channels to differentiate them from any-source multicast (ASM) groups. SSM is ideal for one-to-many multicast services such as network entertainment channels. Although ASM supports one-to-many, its method of source discovery is less efficient than SSM. For example, if you click a link in a browser, ASM notifies the receiver about the group information, but not the source information. With SSM, the client receives both source and group information.

To deploy SSM successfully, you need an end-to-end multicast-enabled network and applications that use an Internet Group Management Protocol version 3 (IGMPv3) or Multicast Listener Discovery version 2 (MLDv2) stack, or you need to configure SSM mapping from IGMPv1 or IGMPv2 to IGMPv3. An IGMPv3 stack provides the capability of a host operating system to use the IGMPv3 protocol.

You can accommodate hosts that do not support IGMPv3 or MLDv1 by using SSM mapping. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report, and MLDv1 reports to MLDv2. SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232/8 for IPv4, ff30::/32 through ff3F::/32 for IPv6).

BEST PRACTICE: Create separate SSM maps for the IPv4 and IPv6 address families when both families require SSM support.

If you apply an SSM map policy containing both IPv4 and IPv6 addresses to an interface in an IPv4 context (using IGMP), only the IPv4 addresses in the list are used. If there are no such addresses, no action is taken. Similarly, if you apply an SSM map policy containing both IPv4 and IPv6 addresses to an interface in an IPv6 context (using MLD), only the IPv6 addresses in the list are used. If there are no such addresses, no action is taken.

To configure SSM mapping for dynamic IGMP:

1. Create an SSM policy to match the desired IPv4, IPv6, or both group addresses.

```
[edit]
user@host# edit policy-options policy-statement policy-name
```

2. Configure terms for the policy to identify and accept group addresses

```
[edit policy-options policy-statement policy-name]
user@host# set term from name route-filter destination-prefix match-type
```

```
user@host# set term name then accept
```

3. Apply the SSM map policy to the dynamic interface in a dynamic profile.

```
[edit dynamic-profiles profile-name protocols (igmp | mld) interface $junos-interface-name]
user@host# set ssm-map-policy ssm-map-policy-name
```

For example, the following configuration creates SSM policy ssm-1. The policy term v4 exactly matches the IPv4 SSM group address 233.252.1.1/32. The policy rejects all other addresses. The policy ssm-1 is then applied to dynamic interfaces created when the igmp-prof dynamic profile is instantiated.

```
[edit]
user@host# edit policy-options policy-statement ssm-1
user@host# set term v4 from route-filter 233.252.1.1/32 exact
user@host# set term v4 then accept
user@host# set then reject
user@host# edit dynamic-profiles mld-prof protocols igmp interface $junos-interface-name
user@host# set ssm-map-policy ssm-1
```

For example, the following configuration creates SSM policy ssm-2. Policy term v6 exactly matches the IPv6 group address ff35::1/128. The policy rejects all other addresses. The policy ssm-2 is then applied to dynamic interfaces created when the mld-prof dynamic profile is instantiated.

```
[edit]
user@host# edit policy-options policy-statement ssm-2
user@host# set term v6 from route-filter ff35::1/128 exact
user@host# set term v6 then accept
user@host# set then reject
user@host# edit dynamic-profiles igmp-prof protocols mld interface $junos-interface-name
user@host# set ssm-map-policy ssm-2
```

RELATED DOCUMENTATION

[Dynamic IGMP Configuration Overview | 393](#)

[Dynamic MLD Configuration Overview | 401](#)

Configuring Dynamic MLD to Enable Subscribers to Access Multicast Networks

IN THIS CHAPTER

- [Dynamic MLD Configuration Overview | 401](#)

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.

Dynamic MLD consists of a subset of the full range of MLD capabilities available for static MLD configuration, applied to dynamic interfaces by means of a dynamic profile. For detailed information about static MLD configuration, see *Configuring MLD*. Much of the static configuration documentation is directly applicable to dynamic MLD. Note that the following statements that appear in the dynamic MLD CLI hierarchy are configurable, but have no effect: **accounting**, **group-threshold**, **log-interval**, and **no-accounting**. These statements are not needed at a subscriber level, where typically no more than tens of joins are expected.

Refer to the *Multicast Protocols User Guide* for a comprehensive understanding of Junos OS support for multicast protocols.

RELATED DOCUMENTATION

Dynamic Profiles Overview

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Configuring MLD

5

PART

Configuring Application-Aware Policy Control and Reporting

[Configuring Application-Aware Policy Control | 405](#)

[Configuring Application Identification | 431](#)

[Configuring Reporting for Application-Aware Data Sessions | 441](#)

Configuring Application-Aware Policy Control

IN THIS CHAPTER

- Understanding Application-Aware Policy Control for Subscriber Management | 406
- Understanding PCC Rules for Subscriber Management | 407
- Configuring Application-Aware Policy Control for Subscriber Management | 410
- Installing Services Packages for Subscriber Management Application-Aware Policy Management | 411
- Configuring Service Data Flow Filters | 412
- Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416
- Configuring Policy and Charging Control Rules | 418
- Configuring a Policy and Charging Control Rulebase | 421
- Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management | 423
- Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control | 424
- Configuring PCC Rule Activation in a Subscriber Management Dynamic Profile | 426
- Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management | 428

Understanding Application-Aware Policy Control for Subscriber Management

Starting in Junos OS Release 16.1R4 and in Junos OS Release 17.2R1, you can configure application-aware policy control, which defines the treatment to apply to a subscriber's packets based on the specific application being used by the subscriber (for example, Facebook) or based on Layer 3 and Layer 4 service data flow (SDF) information for the IP flow (for example, the source and destination IP addresses). Starting in Junos OS Release 19.3R2, application-aware policy control is also supported if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

You configure application-aware policy control by configuring policy and charging control (PCC) rules, which identify the conditions that must be met (such as the application that the traffic is using) and the action to take on that traffic (such as specifying a maximum bit rate). PCC rules can be activated for a subscriber in one of two ways:

- PCC rule activation control by dynamic profile—The dynamic profile assigned to a subscriber identifies a static PCEF profile, which specifies PCC rules. The dynamic profile indicates whether to activate all the rules in the PCEF profile or just a subset of the rules. The PCEF profile and PCC rule names can be variables in the dynamic profile, and the names are obtained by RADIUS during subscriber authorization.
- PCC rule activation by a policy and charging rules function (PCRF) server—Starting in Junos OS Release 18.2R1, a PCRF can directly activate a PCC rule that is configured on the MX Series router by sending a Rule-Install-Name AVP over the Gx interface to the MX Series router during service activation. The specified PCC rule must be identified in a dynamic PCEF profile. If the Rule-Install-Name is also the name of a dynamic profile, then the rule is ignored and the dynamic profile is used.

Benefits

Application-aware policy control allows highly customizable, differentiated services for subscribers.

Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R2, application-aware policy control is also supported if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.
18.2R1	Starting in Junos OS Release 18.2R1, a PCRF can directly activate a PCC rule that is configured on the MX Series router by sending a Rule-Install-Name AVP over the Gx interface to the MX Series router during service activation.
16.1R4	Starting in Junos OS Release 16.1R4 and in Junos OS Release 17.2R1, you can configure application-aware policy control, which defines the treatment to apply to a subscriber's packets based on the specific application being used by the subscriber (for example, Facebook) or based on Layer 3 and Layer 4 service data flow (SDF) information for the IP flow (for example, the source and destination IP addresses).

RELATED DOCUMENTATION

- [Understanding PCC Rules for Subscriber Management | 407](#)
- [Configuring Application-Aware Policy Control for Subscriber Management | 410](#)
- [Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)
- [Configuring Service Data Flow Filters | 412](#)
- [Configuring Policy and Charging Control Rules | 418](#)
- [Application Identification Overview | 431](#)

Understanding PCC Rules for Subscriber Management

IN THIS SECTION

- [Application Filters | 408](#)
- [Service Data Flow Filters | 408](#)
- [PCC Action Profiles | 409](#)

NOTE: Starting in Junos OS Release 19.3R2, PCC rules are also supported if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

Policy and charging control (PCC) rules define the treatment to apply to subscriber traffic based on the application being used by the subscriber (for example, Facebook) or based on the Layer 3 and Layer 4 service data flow (SDF) information for the IP flow (for example, the source and destination IP addresses). You configure PCC rules, and PCC rules are then activated by either the subscriber's dynamic profile or by a PCRF. PCC rules include the following components:

Application Filters

Applications and application groups are specified in the **from** clause of a PCC rule to identify IP packets belonging to a specific application. If the IP packet is for an application identified in a PCC rule, the treatment specified in the PCC action profile in the **then** clause of the rule is applied.

To configure application-aware PCC rules, you can specify one or more of the following parameters:

- **application**—Specifies the name of an application. This can be a Layer 7 protocol (for example, HTTP) or a particular application running on a Layer 7 protocol, such as Facebook and Yahoo Messenger.
- **application-group**—Specifies the name of an application group, which represents a collection of Layer 7 applications that can be processed at the same time.

NOTE: Application-aware PCC rules that reference specified applications can include wildcard or specific Layer-3 SDF filters, Layer-4 SDF filters, or both.

You can see a list of all the applications and application groups by using the **show services application-identification application** command. To configure a custom application, see [“Configuring Custom Application Signatures” on page 435](#).

Service Data Flow Filters

SDF filters (flow identifiers) are specified in the **from** clause of a PCC rule to identify IP packets belonging to a particular Layer 3 or Layer 4 service data flow. If the IP packet matches the SDF filter in a PCC rule, the treatment specified in the PCC action profile in the **then** clause of the rule is applied.

To configure Layer 3 or Layer 4 SDF filters, you specify one or more of the following parameters:

- Source IP address
- Destination IP address

- Source port
- Destination port
- Layer 4 protocol (UDP or TCP)

PCC Action Profiles

A PCC rule configuration includes an action profile in the **then** clause that defines the treatment to apply to a packet belonging to an application or to an SDF identified in the **from** clause of the rule. You can configure a PCC action profile that is used in one or more PCC rules to provide the following functionality:

- HTTP redirection—Specifies HTTP redirection to a URL. You can use this action only for PCC rules that match only HTTP-based applications and all flows.
- HTTP Steering path—Specifies an IPv4 or IPv6 address for steering HTTP packets. You can use this action only for PCC rules that match only HTTP-based applications and all flows.

NOTE: A single PCC rule can support either HTTP redirection or HTTP steering path, but not both.

- Steering with a routing instance—Specifies a routing instance for steering of packets. You can configure different routing instances for traffic from the subscriber (uplink) and traffic to the subscriber (downlink).
- Forwarding class—Specifies the forwarding class that you want assigned to the packet.
- Maximum bit rate—Specifies the maximum bit rate for uplink and for downlink traffic.
- Gating status—Specifies whether to block or to forward IP packets.

RELATED DOCUMENTATION

[Configuring Application-Aware Policy Control for Subscriber Management | 410](#)

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

[Configuring Service Data Flow Filters | 412](#)

[Configuring Policy and Charging Control Rules | 418](#)

Configuring Application-Aware Policy Control for Subscriber Management

This topic gives an overview of the tasks you perform to configure policy control for subscriber management based on the layer 7 application that traffic is using or on the particular Layer 3 or Layer 4 service data flow.

NOTE: Starting in Junos OS Release 19.3R2, application-aware policy control is also supported if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

To configure policy control:

1. Install service packages on any MS-MPC PICs that perform application-aware policy control, or on the MX-SPC3 services card if you have enabled Next Gen Services on either the MX240, MX480, or MX860.

See [“Installing Services Packages for Subscriber Management Application-Aware Policy Management” on page 411.](#)

2. Configure any service data flow filters to be used in PCC rules.

See [“Configuring Service Data Flow Filters” on page 412.](#)

3. Configure any custom applications to be used in PCC rules.

See [“Configuring Custom Application Signatures” on page 435.](#)

4. Configure the PCC action profiles to be used in PCC rules.

See [“Configuring Policy and Charging Control Action Profiles for Subscriber Management” on page 416](#)

5. Configure PCC rules.

See [“Configuring Policy and Charging Control Rules” on page 418.](#)

6. (Optional) Configure PCC rulebases.

See [“Configuring a Policy and Charging Control Rulebase” on page 421.](#)

7. Configure a policy and charging enforcement function (PCEF) profile.

See [“Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management” on page 423.](#)

8. Configure a service set for application-aware policy control.

See [“Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control”](#) on page 424.

9. Perform one of the following:

- For PCC rule activation through a dynamic profile, perform [“Configuring PCC Rule Activation in a Subscriber Management Dynamic Profile”](#) on page 426.
- For direct PCC rule activation by a policy and charging rules function (PCRF) server, perform [“Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management”](#) on page 428.

RELATED DOCUMENTATION

[Understanding Application-Aware Policy Control for Subscriber Management](#) | 406

Installing Services Packages for Subscriber Management Application-Aware Policy Management

You must install a set of service packages on any MS-MPC PICs that perform application-aware policy control, or on the MX-SPC3 services card if you have enabled Next Gen Services on the MX240, MX480, or MX960.

To install service packages:

1. Specify the MS-MPC PIC or MX-SPC3 services card.

```
[edit chassis]
user@host# edit fpc slot-number pic pic-number
```

2. Install the services packages.

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider ]
user@host# set package jservices-mss
user@host# set package jservices-jdpi
user@host# set package jservices-pcef
```

RELATED DOCUMENTATION

Configuring Service Data Flow Filters

NOTE: Starting in Junos OS Release 19.3R2, PCC rules are also supported if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

A service data flow (SDF) filter is specified as a matching condition in the **from** clause of a policy and charging control (PCC) rule. Each SDF filter can have one or more flows associated with it; each flow is a five-tuple match.

NOTE: If you configure an SDF filter without specifying a remote address, port, port range, or protocol, then the SDF filter matches IP packets that have any value configured for the corresponding attribute. If you configure an SDF filter, you must configure at least one of the following attributes: direction, local port or local port range, protocol, remote address, or remote port or remote port range.

You can configure SDF filters for Junos OS Subscriber Aware or for Junos OS Broadband Subscriber Management, but you use a different CLI hierarchy level for each product.

- If you are using Junos OS Subscriber Aware, configure SDF filters at the **[edit unified-edge pcef]** hierarchy level.
- If you are using Junos OS Broadband Subscriber Management, configure SDF filters at the **[edit services pcef]** hierarchy level.

To configure Layer 3 and Layer 4 SDF filters:

1. Specify a name for the SDF filter.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef]  
user@host# set flow-descriptions flow-identifier
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef]  
user@host# set flow-descriptions flow-identifier
```

2. Specify the flow direction for the SDF filter.

NOTE: If you do not specify a flow direction, then the SDF filter is applied in both the uplink and downlink directions.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set direction (uplink | downlink | both)
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set direction (uplink | downlink | both)
```

3. Specify a remote address (IPv4 or IPv6) for the SDF filter:

NOTE: You can specify an IPv4 subnet or an IPv6 subnet but not both.

- Specify an IPv4 address for the SDF filter.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set remote-address ipv4-address ipv4-address
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set remote-address ipv4-address ipv4-address
```

- Specify an IPv6 address for the SDF filter.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set remote-address ipv6-address ipv6-address
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set remote-address ipv6-address ipv6-address
```

4. Specify a protocol (using the standard protocol number) for the SDF filter.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set protocol number
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set protocol number
```

5. Specify a local port or a list of port numbers for the SDF filter. To specify a list of port numbers (up to a maximum of three), enclose the port numbers in square brackets ([]).

NOTE: You can configure a local port or local port range but not both in the same SDF filter.

For Junos OS Subscriber Aware:

```
edit unified-edge pcef flow-descriptions flow-identifier
user@host# set local-ports number
```

For Junos OS Broadband Subscriber Management:

```
edit services pcef flow-descriptions flow-identifier
user@host# set local-ports number
```

6. Specify a local port range for the SDF filter.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set local-port-range low low-value high high-value
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set local-port-range low low-value high high-value
```

7. Specify a remote port or list of remote ports for the SDF filter. To specify a list of port numbers (up to a maximum of three), enclose the port numbers in square brackets ([]).

NOTE: You can configure a remote port or remote port range but not both in the same SDF filter.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set remote-ports number
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set remote-ports number
```

8. Specify a remote port range for the SDF filter.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef flow-descriptions flow-identifier]
user@host# set remote-port-range low low-value high high-value
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef flow-descriptions flow-identifier]
user@host# set remote-port-range low low-value high high-value
```

RELATED DOCUMENTATION

[Configuring Application-Aware Policy Control for Subscriber Management | 410](#)

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

[Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment](#)

[Configuring Policy and Charging Control Rules | 418](#)

Configuring Policy and Charging Control Action Profiles for Subscriber Management

A PCC action profile defines the treatment to be applied to a subscriber's packets associated with specific applications or with specific service data flows. A PCC action profile is specified in the **then** clause of a PCC rule.

NOTE: You cannot change a PCC action profile while it is being used by a subscriber. To modify the PCC action profile, you must log off the subscribers that are using the PCC rule that includes the profile.

To configure PCC action profiles:

1. Specify a name for the PCC action profile.

```
[edit services pcef]
user@host# edit pcc-action-profiles profile-name
```

2. Configure the maximum bit rate for uplink and downlink subscriber traffic.

```
[edit services pcef pcc-action-profiles profile-name]
user@host# set maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value
```

The range is 0 through 6144000 Kbps.

3. Configure HTTP redirection to a URL.

```
[edit services pcef pcc-action-profiles profile-name redirect]
user@host# set url url-name
```

NOTE: A PCC action profile that includes HTTP redirection can only be used in PCC rules that match only HTTP-based applications and all flows.

4. Configure the steering of traffic to a third-party server for applying services or to a service chain with one of the following methods:
 - Specify the IP address of the third-party server for HTTP traffic.

```
[edit services pcef pcc-action-profiles profile-name ]
user@host# set steering path (ipv4-address ipv4-address | set ipv6-address ipv6-address)
```

NOTE: A PCC action profile that includes a steering path can only be used in PCC rules that match only HTTP-based applications and all flows.

- Specify the routing instance to use to reach the third-party server or service chain.

```
[edit services pcef pcc-action-profiles profile-name]
user@host# set steering routing-instance downlink downlink-vrf-name uplink uplink-vrf-name
```

The downlink routing instance is applied to traffic going to the access side, and the uplink routing instance is applied to traffic being sent from the access side.

5. Specify that the PCC action profile steering attributes that a PCC rule applies at the start of a data flow will continue to be applied to that data flow when the PCC rule match conditions are modified, deleted, or added to.

```
[edit services pcef pcc-action-profiles profile-name steering]
user@host# set keep-existing-steering
```

6. Specify the forwarding class to assign to packets.

```
[edit services pcef pcc-action-profiles profile-name]
user@host# set forwarding-class class-name
```

7. Configure the gating status by enabling or disabling the forwarding of packets.

```
[edit services pcef pcc-action-profiles profile-name]
user@host# set gate-status (disable-both | downlink | uplink | uplink-downlink)
```

RELATED DOCUMENTATION

[Configuring Application-Aware Policy Control for Subscriber Management | 410](#)

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

[Configuring Policy and Charging Control Rules | 418](#)

Configuring Policy and Charging Control Rules

A policy and charging control (PCC) rule defines the treatment to be applied to packets associated with specific applications or to specific service data flows.

You can configure PCC rules for Junos OS Subscriber Aware or for Junos OS Broadband Subscriber Management, but you use a different CLI hierarchy level for each product.

- If you are using Junos OS Subscriber Aware, configure PCC rules at the **[edit unified-edge pcef]** hierarchy level.
- If you are using Junos OS Broadband Subscriber Management, configure PCC rules at the **[edit services pcef]** hierarchy level.

NOTE: If you are using Junos OS Subscriber Aware, you must be in maintenance mode to make a change to a PCC rule. (See *Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles*).

NOTE: If you are using Junos OS Broadband Subscriber Management, you cannot change a PCC rule while it is being used by a subscriber. To modify the rule, you must log off the subscribers that are using the rule.

Before you configure PCC rules, you must do the following:

- Configure the service data flow (SDF) filters that the PCC rules reference.
- Configure the application groups and any custom applications that you want to reference in application-aware PCC rules.
- Configure the PCC action profiles that the PCC rules reference.

NOTE: When specifying application-aware PCC rules in a PCEF profile, you must also configure a default Layer 3 or Layer 4 wildcard PCC rule to ensure that the default charging characteristics are applied to unmatched subscriber traffic without dropping that traffic. For example, the default Layer 3 or Layer 4 wildcard PCC rule prevents traffic based on DNS queries from being dropped. In addition, the policy (PCEF profile) that includes application-aware PCC rules must also include a wildcard Layer 3 or Layer 4 PCC rule at a lower precedence.

To configure PCC rules:

1. Specify a name for the PCC rule.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef]
user@host# edit pcc-rules rule-name
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef]
user@host# edit pcc-rules rule-name
```

2. In a **from** statement, specify an SDF filter to use Layer 3 or Layer 4 match conditions for filtering subscriber traffic.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rules rule-name]
user@host# set from flows flow-identifier
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rules rule-name]
user@host# set from flows flow-identifier
```

If you do not want to filter subscriber traffic based on SDF filters, use the **any** option.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rules rule-name]
user@host# set from flows any
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rules rule-name]
user@host# set from flows any
```

3. (Optional) Specify an application as a match condition for filtering subscriber traffic.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rules rule-name]
user@host# set from applications application-name
```

For Junos OS Broadband Subscriber Management:


```
[edit services pcef pcc-rules rule-name]
user@host# set from applications application-name
```

4. (Optional) Specify multiple applications instead of specifying each application separately by specifying an application group as a match condition for filtering subscriber traffic.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rules rule-name]
user@host# set from application-groups application-group-name
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rules rule-name]
user@host# set from application-groups application-group-name
```

5. Specify the PCC rules action profile that defines the treatment to be applied to specific service data flows or to packets associated with specific applications.

NOTE: You can use PCC action profiles with HTTP redirection or HCM profiles only in PCC rules that match only HTTP-based applications and any flows.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rules rule-name]
user@host# set then pcc-action-profile profile-name
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rules rule-name]
user@host# set then pcc-action-profile profile-name
```

RELATED DOCUMENTATION

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Understanding Application-Aware Policy Control for Subscriber Management](#) | 406

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416

Configuring Service Data Flow Filters | 412

Configuring Custom Application Signatures | 435

Configuring a Policy and Charging Control Rulebase

A policy and charging control (PCC) rulebase contains a set of PCC rules. Each rule specified in the PCC rulebase is assigned a precedence to designate the priority in which PCC rules are evaluated for selection in a policy and charging enforcement function (PCEF) profile.

NOTE: Starting in Junos OS Release 19.3R1, application-aware policy control is also supported for Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

You can configure PCC rulebases for Junos OS Subscriber Aware or for Junos OS Broadband Subscriber Management, but you use a different CLI hierarchy level for each product.

- If you are using Junos OS Subscriber Aware, configure PCC rulebases at the **[edit unified-edge pcef]** hierarchy level.
- If you are using Junos OS Broadband Subscriber Management, configure PCC rulebases at the **[edit services pcef]** hierarchy level.

NOTE: If you are using Junos OS Subscriber Aware, you must be in maintenance mode to make a change to a PCC rulebase. (See *Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles*).

NOTE: If you are using Junos OS Broadband Subscriber Management, you cannot change a PCC rulebase while it is being used by a subscriber. To modify the rulebase, you must log off the subscribers that are using the rule.

Before you configure a PCC rulebase, you must do the following:

- Configure service data flow filters.
- Configure PCC action profiles.
- Configure PCC rules.

To configure a PCC rulebase:

1. Specify a name for the rulebase.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef ]
user@host# edit pcc-rulebases rulebase-name
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef ]
user@host# edit pcc-rulebases rulebase-name
```

2. Specify the PCC rules that the rulebase references and a precedence value (1 through 65,535) for each rule.

NOTE:

- The same rule can be configured in different rulebases and can have a different precedence.
- The precedence assigned must be unique among the configured PCC rules.
- A lower precedence value indicates a higher precedence. For example, if a PCC rulebase has two PCC rules with precedence 5 and 10 respectively, the PCC rule with precedence 5 is evaluated first.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-rulebases rulebase-name]
user@host# set pcc-rule rule-name precedence number
user@host# set pcc-rule rule-name precedence number
user@host# set pcc-rule rule-name precedence number
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-rulebases rulebase-name]
user@host# set pcc-rule rule-name precedence number
user@host# set pcc-rule rule-name precedence number
user@host# set pcc-rule rule-name precedence number
```

RELATED DOCUMENTATION

[Configuring Policy and Charging Control Rules | 418](#)

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

Understanding Junos Subscriber Aware Policy and Charging Enforcement Function (PCEF)

Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management

NOTE: Starting in Junos OS Release 19.3R2, policy and charging enforcement function (PCEF) profiles are also supported if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

A PCEF profile specifies a set of PCC rules and rulebases that can be assigned to a subscriber, and assigns a precedence value to each predefined rule. The PCEF profile is used in one of the following ways:

- A static PCEF profile is specified in a dynamic profile. The dynamic profile indicates whether to activate all the rules in the PCEF profile or just a subset of the rules.
- A dynamic PCEF profile identifies the PCC rules and rulebases that a PCRF can directly activate.

NOTE: You cannot change a PCEF profile while it is being used by a subscriber. To modify the PCEF profile, you must log off the subscribers that are using the PCEF profile.

To configure a PCEF profile:

1. Specify a name for the PCEF profile.

```
[edit services pcef]
user@host# edit profiles profile-name
```

2. Specify one or more PCC rules and a precedence for each rule. A lower precedence value indicates a higher precedence. The precedence assigned must be unique among the configured PCC rules, including the PCC rules that are assigned a precedence within a PCC rulebase.
 - For a PCEF profile that is specified in a dynamic profile, specify the rules under static-policy-control.

```
[edit services pcef profiles profile-name]
user@host# set static-policy-control pcc-rules rule-name precedence number
```

- For a PCEF profile that identifies the PCC rules that a PCRF can directly activate, specify the rules under dynamic-policy-control.

```
[edit services pcef profiles profile-name]
user@host# set dynamic-policy-control pcc-rules rule-name precedence number
```

3. Specify one or more PCC rulebases.

- For a PCEF profile that is specified in a dynamic profile, specify the rulebases under static-policy-control.

```
[edit services pcef profiles profile-name]
user@host# set static-policy-control pcc-rulebases rulebase-name
```

- For a PCEF profile that identifies the PCC rules that a PCRF can directly activate, specify the rulebases under dynamic-policy-control.

```
[edit services pcef profiles profile-name]
user@host# set dynamic-policy-control pcc-rulebases rulebase-name
```

RELATED DOCUMENTATION

[Configuring Application-Aware Policy Control for Subscriber Management | 410](#)

[Configuring Policy and Charging Control Rules | 418](#)

[Configuring a Policy and Charging Control Rulebase | 421](#)

Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control

NOTE: Starting in Junos OS Release 19.3R2, application-aware policy control is also supported if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

Configure a service set to identify the service interface that handles application-aware policy control.

To configure a service set for application-aware policy control:

1. Define an application-aware service set.

```
[edit services]
user@host# set service-set service-set-name service-set-options subscriber-awareness
```

2. Enable PCEF services for the service set by specifying a dummy name for the **pcef-profile**.

- a. Configure a dummy PCEF profile.

```
[edit services pcef]
user@host# set profiles profile-name
```

- b. Specify the dummy profile in the service set.

```
[edit services service-set service-set-name]
user@host# set pcef-profile pcef-profile-name
```

3. Enable application identification for the service set by specifying a dummy name for the **application-identification-profile**.

- a. Configure a dummy application identification profile.

```
[edit services application-identification]
user@host# set profile app-id-profile-name
```

- b. Specify the dummy profile in the service set.

```
[edit services service-set service-set-name]
user@host# set application-identification-profile app-id-profile-name
```

4. Specify the services PIC interface on which the services are performed.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

If you have redundancy configured, the *interface-name* is *amsn* if you do not have Next Gen Services enabled, and is *ams0.1* if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

If you do not have redundancy configured, the *interface-name* is *ms-fpc/pci/0* if you do not have Next Gen Services enabled on the MX-SPC3 services card on the MX240, MX480, or MX860 router, and is *vsp-fpc/pci/0* if you do have Next Gen Services enabled for MX-SPC3 services card on the MX router.

RELATED DOCUMENTATION

[Configuring Application-Aware Policy Control for Subscriber Management | 410](#)

Configuring PCC Rule Activation in a Subscriber Management Dynamic Profile

NOTE: Starting in Junos OS Release 19.3R2, PCC rule activation in a dynamic profile is also supported if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

To configure PCC rule activation by a dynamic profile, specify the PCEF profile to use, the PCC rules to activate, and the service set to use.

1. Assign a PCEF profile to the dynamic profile. In the client dynamic profile, you can identify the PCEF profile with the variable **\$junos-pcef-profile**. All of a subscriber's dynamic profiles that include a PCEF profile must point to the same PCEF profile.

```
[edit dynamic-profile profile-name interfaces interface-name unit logical-unit-number service]
user@host# set pcef pcef-profile-name
```

2. Activate PCC rules in the dynamic profile. In the access profile, you can identify a rule name with the variable **\$junos-pcef-rule**.

NOTE: Do not activate both service data flow (Layer 3 or Layer 4) PCC rules that have a gating action and application-aware (Layer 7) PCC rules in the same dynamic profile. The gating action for the service data flow PCC rules is not applied in this situation.

To activate one PCC rule:

```
[edit dynamic-profile profile-name interfaces interface-name unit logical-unit-number service pcef
  pcef-profile-name]
user@host# set activate rule-name
```

To activate all the PCC rules:

```
[edit dynamic-profile profile-name interfaces interface-name unit logical-unit-number service pcef
  pcef-profile-name]
user@host# set activate-all
```

3. Assign a service set to the dynamic profile. This must be a service set that you configured for application-aware policy control. In the client dynamic profile, you can identify the service set with a variable (**\$junos-input-service-set** | **\$junos-output-service-set** | **\$junos-input-ipv6-service-set** | **\$junos-output-ipv6-service-set**). You must use the same service set for both the input and output service.

```
[edit dynamic-profile profile-name interfaces interface-name unit logical-unit-number family family service (input
  | output) service-set]
user@host# set service-set service-set-name
```

4. (Optional) Assign a service filter to the dynamic profile. The service filter can identify conditions for which you want to skip application-aware policy control. In the client dynamic profile, you can identify the service filter with a variable (**\$junos-input-service-filter** | **\$junos-output-service-filter** | **\$junos-input-ipv6-service-filter** | **\$junos-output-ipv6-service-filter**).

```
[edit dynamic-profile profile-name interfaces interface-name unit logical-unit-number family family service (input
  | output) service-set service-set-name]
user@host# set service-filter filter-name
```

RELATED DOCUMENTATION

[Configuring Application-Aware Policy Control for Subscriber Management](#) | 410

[Understanding Application-Aware Policy Control for Subscriber Management](#) | 406

Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management

NOTE: Starting in Junos OS Release 19.3R1, direct PCC rule activation by a PCRF is also supported if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

Enable direct PCRF activation of PCC rules by configuring a PCRF partition, a Diameter instance, and a PCC context in an access profile.

1. Configure the Diameter instance. See *Configuring Diameter*.
2. Configure the PCRF partition. See *Configuring the PCRF Partition*.
3. Enable PCRF provisioning in the access profile.

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

4. Assign a PCEF profile to the access profile PCC context.

```
[edit access profile profile-name session-options pcc-context]
user@host# set profile-name pcef-profile-name
```

5. Specify the IPv6 input service set that handles application-aware policy control.

```
[edit access profile profile-name session-options pcc-context]
user@host# set ipv6-input-service-set-name service-set-name
```

6. (Optional) Specify a service filter for the IPv6 input service set to identify conditions for which you want to skip application-aware policy control.

```
[edit access profile profile-name session-options pcc-context]
user@host# set ipv6-input-service-filter-name filter-name
```

7. Specify the IPv4 input service set that handles application-aware policy control.

```
[edit access profile profile-name session-options pcc-context]
```

```
user@host# set input-service-set-name service-set-name
```

8. (Optional) Specify a service filter for the IPv4 input service set to identify conditions for which you want to skip application-aware policy control.

```
[edit access profile profile-name session-options pcc-context]
user@host# set input-service-filter-name filter-name
```

9. Specify the IPv6 output service set that handles application-aware policy control.

```
[edit access profile profile-name session-options pcc-context]
user@host# set ipv6-output-service-set-name service-set-name
```

10. (Optional) Specify a service filter for the IPv6 output service set to identify conditions for which you want to skip application-aware policy control.

```
[edit access profile profile-name session-options pcc-context]
user@host# set ipv6-output-service-filter-name filter-name
```

11. Specify the IPv4 output service set that handles application-aware policy control.

```
[edit access profile profile-name session-options pcc-context]
user@host# set output-service-set-name service-set-name
```

12. (Optional) Specify a service filter for the IPv4 output service set to identify conditions for which you want to skip application-aware policy control.

```
[edit access profile profile-name session-options pcc-context]
user@host# set output-service-filter-name filter-name
```

RELATED DOCUMENTATION

[Configuring Application-Aware Policy Control for Subscriber Management | 410](#)

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

Configuring Application Identification

IN THIS CHAPTER

- [Application Identification Overview | 431](#)
- [Downloading and Installing Predefined Junos OS Application Signature Packages | 432](#)
- [Improving the Application Traffic Throughput | 434](#)
- [Configuring Custom Application Signatures | 435](#)
- [Uninstalling a Predefined Junos OS Application Signature Package | 440](#)

Application Identification Overview

Junos Application Aware is an infrastructure plug-in on MS-MPC service PICs and on the MX-SPC3 services card that provides information to clients about application protocol bundles based on deep packet inspection (DPI) of application signatures. These clients can be any of the plug-ins on the MX Series router service chain, such as traffic detection function (TDF), that request application classification data. Starting in Junos OS Release 16.1R4 and Junos OS Release 17.2R1, application identification is available in Junos OS Broadband Subscriber Management. Starting in Junos OS Release 19.3R2, application identification is also supported for Broadband Subscriber Management on the MX-SPC3 services card if you have enabled Next Gen Services on the MX240, MX480 or MX960 router.

In application identification, you can apply application signatures as follows:

- **Predefined signatures**—Junos Application Aware comes with a bundle of predefined, preinstalled application signatures, but we recommend that you download and install the latest version of predefined signatures. As new sets of signatures are supported, they are compiled and made available for you to download.
- **Custom application signatures**—For any application signatures that are not predefined, you can create custom signatures for HTTP, SSL, and stream signature contexts and install them for application identification. After you have configured and committed custom signatures, they are serialized and merged with the predefined application signatures. You can specify the following types of custom application signatures:

- **Address based**—You can define an application identification based on a specific IP address, or port, or both where a source IP address, destination IP address, or both are used for a known application in a customer's network. This is useful, for example, when a Session Initiation Protocol (SIP) server initiates a session from its well known port, 5060. The customer can put the SIP server IP address and port 5060 as source IP/port for the SIP application. This method provides efficiency and accuracy of application identification for customer's network.
- **Internet Control Message Protocol (ICMP) based**—Application identification based on types of ICMP messages.
- **IP protocol based**—Application identification based on IP protocol. TCP, UDP, and ICMP are not supported for this method of signature creation.
- **Pattern-matching signatures**—Application based on pattern matching combined with Layer 7 protocol identification.

Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R2, application identification is also supported for Broadband Subscriber Management on the MX-SPC3 services card if you have enabled Next Gen Services on the MX240, MX480 or MX960 router.
16.1R4	Starting in Junos OS Release 16.1R4 and Junos OS Release 17.2R1, application identification is available in Junos OS Broadband Subscriber Management.

RELATED DOCUMENTATION

[Configuring Custom Application Signatures | 435](#)

[Downloading and Installing Predefined Junos OS Application Signature Packages | 432](#)

Downloading and Installing Predefined Junos OS Application Signature Packages

NOTE: Starting in Junos OS Release 19.3R1, application identification is also supported for Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card..

To download, install, and verify the installation of predefined Junos OS application signature packages:

1. Use **download ignore-server-validation** if you want to skip server certification validation during the download. Validation is enabled by default.

```
[edit services application-identification]
user@host# set download ignore-server-validation
```

2. Configure the URL for the application signature packages server.

```
[edit services application-identification]
user@host# set download url https://services.netscreen.com/cgi-bin/index.cgi
```

3. Download the application signature package.

- To download the latest signature package, enter the following command:

```
user@host> request services application-identification download
```

- To download a specific, known signature package, include the version number:

```
user@host> request services application-identification download version version-number
```

4. Confirm the successful download of the package.

```
user@host> request services application-identification download status
```

```
Downloading application package succeed.
```

5. Install the application signature package.

```
user@host> request services application-identification install
```

6. Confirm the successful installation of the application signature package.

```
user@host> request services application-identification install status
```

```
Compiling application signatures of package version.
```

or

```
Install application package succeed
```

7. View the protocol bundle status:

```
user@host> show services application-identification status
```

RELATED DOCUMENTATION

[Uninstalling a Predefined Junos OS Application Signature Package | 440](#)

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

Improving the Application Traffic Throughput

NOTE: Starting in Junos OS Release 19.3R1, application identification is also supported for Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

To improve the application throughput when using application identification for deep packet inspection (DPI):

1. Enable the DPI performance mode.

```
[edit services application-identification]  
user@host# set enable-performance-mode
```

This limits the maximum DPI processing to four packets per session.

By default, DPI performance mode is disabled.

Configuring Custom Application Signatures

NOTE: Starting in Junos OS Release 19.3R1, application identification is also supported for Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

You can configure custom application definitions using custom signatures. These definitions enable identification of protocol bundles through deep packet inspection (DPI) for use by interested services in the service chain.

Before you configure custom application signatures, ensure that **jservices-jdpi** is configured on all required interfaces of your MS-MPC, or of your MX-SPC3 services card if you have enabled Next Gen Services on the MX240, MX480, or MX960. To review how to configure the package on your MS-MPC or MX-SPC3 services card:

- For Junos OS Subscriber Aware, see *Preconfigured Groups for Service PICs and for Session PICs Overview*.
- For Junos OS Broadband Subscriber Management, see [“Installing Services Packages for Subscriber Management Application-Aware Policy Management” on page 411](#).

To configure one or more custom application signatures:

1. Specify a name for the application.

```
[edit services application-identification]
user@host# edit application application-name
```

For example:

```
[edit services application-identification]
user@host# edit application my:http
```

2. Specify a description for the application.

```
[edit services application-identification application application-name]
user@host# set description description
```

For example:

```
[edit services application-identification application my:http]
user@host# set description "Test application"
```

3. Specify an alternative name for the application.

```
[edit services application-identification application application-name]
user@host# set alt-name alt-name
```

For example:

```
[edit services application-identification application my:http]
user@host# set alt-name my:http-app
```

4. Enable saving of the application system cache (ASC).

```
[edit services application-identification application my:http]
user@host# set cacheable
```

5. Specify the name of the Junos OS release for compatibility.

```
[edit services application-identification application application-name]
user@host# set compatibility junos-compatibility-version
```

For example:

```
[edit services application-identification application my:http]
user@host# set compatibility 17.1
```

6. Specify any desired application tags, consisting of a user-defined name and value.

```
[edit services application-identification application application-name]
user@host# set tags tag-name tag-value
```

For example:

```
[edit services application-identification application my:http]
user@host# set tags traffic-type video-stream
```

7. Specify one or more address-based signatures.

- Specify a destination address and destination port-range.

```
[edit services application-identification application application-name]
```



```
user@host# set filter ip 200.0.0.2/24 port-range [80]
```

8. Specify an ICMP-based signature.

- a. Specify ICMP type and code.

```
[edit services application-identification application application-name]
user@host# set icmp-mapping type icmp-type code icmp-code
```

For example:

```
[edit services application-identification application my:http]
user@host# set icmp-mapping type 33 code 34
```

9. Specify an IP protocol-based signature.

- a. Specify the IP protocol by protocol number.

```
[edit services application-identification application application-name]
user@host# set ip-protocol-mapping protocol protocol-number
```

For example:

```
[edit services application-identification application my:http]
user@host# set ip-protocol-mapping protocol 103
```

All ip-protocol-mappings are allowed except Protocol numbers 1,6,17 are not allowed to be configured under ip-protocol based signatures. If you try to configure protocols 1,6,17 under ip-protocol-mapping you will get commit errors.

10. Specify one or more Layer 4 and Layer 7 signatures using pattern matching in conjunction with a Layer 4 protocol.

- a. Specify a name for the Layer 4 and Layer 7 signature.

```
[edit services application-identification application application-name over protocol-type]
user@host# set signature I4-I7-signature-name
```

For example:

```
[edit services application-identification application my:http over http]
user@host# set signature myI3I7
```

- b. Specify the order to be used if conflicts occur during the application classification. In such a case, the application with lowest order is classified.

```
[edit services application-identification application application-name over protocol-type signature
  l4-l7-signature-name member member-name]
user@host# set order order
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7 member m01]
user@host# set order 1
```

- c. Specify the priority for using this signature instead of using any matched predefined signatures.

```
[edit services application-identification application application-name over protocol-type signature
  l4-l7-signature-name]
user@host# set order-priority (high | low)
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7]
user@host# set order-priority high
```

- d. (Optional) Specify the protocol. If you are using Next Gen Services with the MX-SPC3 services card, do not perform this step.

```
[edit services application-identification application application-name over protocol-type signature
  l4-l7-signature-name]
user@host# set protocol (http | ssl | tcp | udp)
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7]
user@host# set protocol http
```

- e. (Optional) Specify that members are to be matched in order.

```
[edit services application-identification application application-name over protocol-type signature
  l4-l7-signature-name]
user@host# set chain-order
```

- f. Specify a member. You can repeat this step to define up to four members.

```
[edit services application-identification application application-name over protocol-type signature
l4-l7-signature-name]
user@host# edit member member-name
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7]
user@host# edit member m01
```

- g. Specify the member's identifying pattern.

```
[edit services application-identification application application-name over protocol-type signature
l4-l7-signature-name member member-name]
user@host# set pattern pattern
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7 member m01]
user@host# set pattern "www.facebook.net"
```

- h. Specify the direction of flows to which pattern matching is applied.

```
[edit services application-identification application application-name over protocol-type signature
l4-l7-signature-name member member-name]
user@host# set direction (any | client-to-server | server-to-client)
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7 member m01]
user@host# set direction any
```

- i. Specify the number of check-bytes. This option applies to TCP and UDP only.

```
[edit services application-identification application application-name over protocol-type signature
l4-l7-signature-name member member-name]
user@host# set check-bytes max-bytes-to-check
```

For example:

```
[edit services application-identification application my:http over http signature myl3l7 member m01]
user@host# set check-bytes 5000
```

11. (For Next Gen Services with the MX-SPC3 services card only) After you have committed your changes, you can check the status of the custom signature commitment.

```
[edit services application-identification application my:http over http signature myl3l7 member m01]  
user@host> show services application-identification commit-status
```

RELATED DOCUMENTATION

| [Application Identification Overview | 431](#)

Uninstalling a Predefined Junos OS Application Signature Package

NOTE: Starting in Junos OS Release 19.3R1, application identification is also supported for Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card.

To uninstall the current application signature package:

- Enter the uninstall command.

```
user@host> request service application-identification uninstall
```

RELATED DOCUMENTATION

| [Downloading and Installing Predefined Junos OS Application Signature Packages | 432](#)

Configuring Reporting for Application-Aware Data Sessions

IN THIS CHAPTER

- [Logging and Reporting Function for Subscribers | 441](#)
- [Log Dictionary for Template Types | 448](#)
- [Configuring Logging and Reporting for Subscriber Management | 459](#)
- [Installing Services Packages for Subscriber Management Logging and Reporting | 459](#)
- [Configuring an LRF Profile for Subscribers | 460](#)
- [Applying Logging and Reporting Configuration to a Subscriber Management Service Set | 467](#)
- [Configuring the Activation of an LRF Rule by a PCC Rule | 468](#)

Logging and Reporting Function for Subscribers

IN THIS SECTION

- [Log and Report Control | 442](#)
- [Templates | 442](#)
- [HTTP Transaction Logging | 447](#)

The logging and reporting function (LRF) enables you to log data for subscriber application-aware policy control sessions and send that data in an IPFIX format to an external log collector using UDP-based transport. These data session logs can include subscriber information, application information, HTTP metadata, data volume, time-of-day information, and source and destination details.

Starting in Junos OS Release 16.1R4 and in Junos OS Release 17.2R1, LRF is available in Junos OS Broadband Subscriber Management. Starting in Junos OS Release 19.3R2, LRF is available in Junos OS Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card..

The external collector, which is not a Juniper Networks product, can then use this data to perform analytics that provide you with insights about subscriber and application usage, allowing you to create packages and policies that increase revenue.

Log and Report Control

A subscriber's data sessions are logged and sent to collectors based on an LRF profile that you configure and associate with the subscriber.

The LRF profile includes:

- **Templates**—Specify the type of data that you want sent and the trigger that causes data to be sent. You can configure a maximum of 16 templates in an LRF profile.
- **Collectors**—Identify the destination to send data to. You can configure a maximum of eight collectors in an LRF profile.
- **LRF rules**—Specify the template and collector to use and, if applicable, a data volume limit that triggers the sending of data. An LRF rule's actions are performed when the matching conditions in a static PCC rule that references the LRF rule are met. You can configure a maximum of 32 LRF rules in an LRF profile.

To associate the LRF profile with a subscriber:

- For Junos OS Subscriber Aware, assign the LRF profile to the subscriber-aware TDF service set that belongs to the TDF interface (mif) in the subscriber's TDF domain.
- For Junos OS Broadband Subscriber Management, assign the LRF profile to the service set that is configured for application-aware policy control.

Templates

NOTE: If you have enabled Next Gen Services with the MX-SPC3 services card, then the DNS, IPv4 extended, IPv6 extended, mobile subscriber, video, and wireline subscriber templates are not supported.

You specify the data fields in a template by configuring one or more types for the template; for example, HTTP and IPv4. Each type represents a set of fields, and the template you configure includes fields from all the types you configure. The template is sent to the collector when you configure it, and is re-sent at a configurable interval. The template types that you can select and the fields that are included by each type are:

- **Device Data**—Contains data fields specific to the device collecting the logging feed:
 - DPI Engine Version
 - IP address of TDF gateway (in IPv4 format)
- **DNS**—(Not available if Next Gen Services is enabled with the MX-SPC3 services card) Contains the DNS response time data field.
- **Flow ID**—Contains the Flow ID data field.

When HTTP multiple transaction logging is enabled, FlowID is an implicit type that gets included with the HTTP template. When the consolidated session log is generated at the time of `SESSION_CLOSE`, LRF includes the FlowID that can be used to correlate with the HTTP transaction log records.

- **HTTP**—Contains data fields for the HTTP metadata from header fields:
 - User Agent
 - Content Length - Request
 - HTTP Response Code
 - Language
 - Host
 - Location
 - Http Method
 - Referer (HTTP)
 - MIME type
 - Time to First Byte
- **IFL subscriber**— Contains data fields specific to IFL-based subscribers:
 - **Subscriber Name**—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).
 - **IFL Name**—Filled with default IFL name (filled with values Next Gen Services IFL)
- **IPFlow**—Contains data fields for the uplink and downlink octets and bytes. When a data record for volume limit is exported, these IPFlow statistics in the record are the actual data received after the last volume limit was reported in that data session and *not* cumulative data.

- Uplink Octets
- Downlink Octets
- Uplink Packets
- Downlink Packets
- Ip Protocol—Protocol ID from IP header; for example, 17 (UDP), 6 (TCP).
- Record Reason—A value of **1** for the session close and a value of **2** for volume-limit.
- IPFlow Extended—Contains data fields for the service set name, routing instance, and payload timestamps. The initiator of the very first packet of a session is the client and the responder is the server.
 - Service-Set-Name—Filled with active **service-set-name** (16 byte value is filled active **service-set-name**. For example, if **service-set-name** is: bng-service-set-1, the template has a value of: bng-service-set-(16bytes)
 - Routing-Instance—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).
- IPFlow TCP—Contains data fields for TCP-related timestamps:
 - Retransmitted TCP packets uplink
 - Retransmitted TCP packets downlink
 - TCP flow creation timestamp
- IPFlow TCP Timestamp—Contains IBM-specific data fields for TCP-related timestamps:
 - Smooth RTT uplink
 - Smooth RTT downlink
 - Client setup time
 - Server Setup time
 - First Client Payload timestamp
 - Upload time
 - First Server Payload timestamp
 - Download time
 - Acknowledged volumes uplink
 - Acknowledged volumes downlink

To use the IPFlow TCP Timestamp template when configuring an LRF profile, identify the template as vendor specific to avoid a commit warning. See [“Configuring an LRF Profile for Subscribers” on page 460](#).

- IPFlow Timestamp—Contains data fields for the flow start and end timestamps:

- Flow Start Time—For TCP, the flow start time is when the SYN packet is received. For UDP, it is when the first packet is sent.
- Flow End Time
- IPv4—Contains data fields for the basic source and destination IPv4 information:
 - Source IPv4 Address
 - Destination IPv4 Address
- IPv4 Extended—(Not available if Next Gen Services with the MX-SPC3 services card are enabled) Contains data fields for the elements of IPv4 extended fields:
 - IPv4 TOS / Class of Service
 - IPv4 Source Mask
 - IPv4 Destination Mask
 - IPv4 Next Hop
- IPv6—Contains data fields for the basic source and destination IPv6 information:
 - Source IPv6 Address
 - Destination IPv6 Address
- IPv6 Extended—(Not available if Next Gen Services are enabled with the MX-SPC3 services card) Contains data fields for the elements of IPv6 extended fields:
 - IPv6 Source Mask
 - IPv6 Destination Mask
 - IPv6 Next Hop
 - Traffic Class
- L7 Application—Contains data fields for the Layer 7 application:
 - Application Protocol—Application data protocol below the classified application name; for example, **http** or **ssl**.
 - Application Name—Application name; for example, **junos:facebook** or **junos:Netflix**.
 - Host—HTTP header host when application protocol is **http**, SSL common name when application protocol is **ssl**, DNS name when application protocol is **dns**.
- Mobile Subscriber—(Not available if Next Gen Services with the MX-SPC3 services card are enabled) Contains data fields specific to mobile subscribers:
 - IMSI
 - MSISDN
 - IMEI

- RAT-type
- ULI
- RADIUS Called Station ID
- PCC—Contains the PCC rule name data field. Not applicable if Next Gen Services are enabled.
- Status Code Distribution—Contains data fields for the HTTP or DNS status codes:
 - Status code 1
 - Status code 2
 - Status code 3
 - Status code 4
 - Status code 5
 - Num Instances 1
 - Num Instances 2
 - Num Instances 3
 - Num Instances 4
 - Num Instances 5
- Subscriber Data—Contains data fields for Generic Subscriber information that can be included with wireless (mobile) subscribers or wireline subscribers:
 - NAS_IP_ADDR—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).
 - Subscriber Type—1 for IP-based subscriber, 2 for IFL-based subscriber.
 - Subscriber IP Address
 - Subscriber VRF—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).
 - NAS Port ID—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).
 - Accounting-Session-Id—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).
 - Class—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).
 - NAS Port Type—Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).
- Transport Layer—Contains data fields for the transport layer:
 - Source Transport Port

- Destination Transport Port
- Video—(Not available if Next Gen Services with the MX-SPC3 services card are enabled) Contains data fields for video traffic:
 - Bitrate
 - Duration
- Wireline Subscriber—(Not available if Next Gen Services with the MX-SPC3 services card are enabled) Contains the UserName data field for wireline subscribers. This is the same as RADIUS Called Station ID.

The template that is specified in an LRF rule determines the set of data fields that are included when data is sent to a collector. The data message includes a pointer to the template ID so that the collector can correlate the data contents with the data field lengths and types.


In a template, you also specify the type of trigger that determines when to send data to the collector. This trigger type can be a data volume limit, a time limit, or the closing of a data session (UDP sessions are considered closed after 60 seconds of inactivity; TCP sessions are considered closed when a FIN, FIN-ACK, or RST is received).

HTTP Transaction Logging

You may enable HTTP transaction logging in an LRF profile. This causes each HTTP transaction in a TCP session to be separately logged and sent to the collector, as shown in [Figure 7 on page 448](#). This option is only relevant when the template being used includes HTTP in the template type.

By default, HTTP transaction logging is disabled, and the HTTP transaction records for a TCP session are sent together as one group of records.

Figure 7: HTTP Transaction Logging



Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R2, LRF is available in Junos OS Broadband Subscriber Management if you have enabled Next Gen Services on the MX240, MX480 or MX960 router with the MX-SPC3 card..
16.1R4	Starting in Junos OS Release 16.1R4 and in Junos OS Release 17.2R1, LRF is available in Junos OS Broadband Subscriber Management.

RELATED DOCUMENTATION

- Log Dictionary for Template Types | 448
- Configuring Logging and Reporting for Junos OS Subscriber Aware
- Configuring Logging and Reporting for Subscriber Management | 459

Log Dictionary for Template Types

Table 36 on page 449 shows the logging dictionary of the template types that LRF supports. The log fields are a mix of IETF standard fields and fields that Juniper Networks defined. The IPFIX convention for vendor-defined fields is an enterprise bit set to 1 and an enterprise ID set to the vendor-ID. (The Juniper

Networks vendor-ID is 2636.) An IETF standard field has an enterprise bit set to **0** and no value for the enterprise ID.

NOTE: If you have enabled Next Gen Services with the MX-SPC3 services card, then the DNS, IPv4 extended, IPv6 extended, mobile subscriber, video, and wireline subscriber templates are not supported.

Table 36: Logging Dictionary for Template Types

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
Device Data	DPI Engine Version	1/2636	503	string	32
	IP address of TDF gateway.	1/2636	502	ipv4Address	4
DNS (Not available if Next Gen Services with the MX-SPC3 services card are enabled)	DNS response time	1/2636	876	dateTimeMilliseconds	8
Flow ID	Flow ID	1/2636	107	unsigned32	4

Table 36: Logging Dictionary for Template Types (continued)

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
HTTP	User Agent	1/2636	152	string	32
	Content Length - Request	1/2636	154	unsigned32	4
	HTTP Response Code	1/2636	155	unsigned16	2
	Language	1/2636	156	string	16
	Host	1/2636	157	string	64
	Location	1/2636	158	string	64
	Http Method	1/2636	159	string	8
	Referer(HTTP)	1/2636	160	string	64
	MIME type	1/2636	161	string	32
	Http URI	1/2636	163	string	255
	Time to First Byte	1/2636	181	dateTimeMilliseconds	8
IFL Subscriber	Subscriber Name	1/2636	511	string Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	16
	IFL Name	1/2636	512	string Filled with default IFL name (filled with values Next Gen Services IFL)	16

Table 36: Logging Dictionary for Template Types (*continued*)

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
IPFlow	Uplink Octets	1/2636	103	unsigned32	4
	Downlink Octets	1/2636	104	unsigned32	4
	Uplink Packets	1/2636	105	unsigned32	4
	Downlink Packets	1/2636	106	unsigned32	4
	Ip Protocol	0	4	unsigned8	1
	Record Reason	1/2636	112	unsigned8	1

Table 36: Logging Dictionary for Template Types (continued)

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
IPFlow Extended	Service-Set-Name	1/2636	520	string Contains data fields for the service-set-name , routing-instance , and payload timestamps. The initiator of the very first packet of a session is the client and the responder is the server. Filled with active service-set-name (16 byte value is filled active service-set-name . For example, if service-set-name is: bng-service-set-1, the template has a value of: bng-service-set-(16bytes)	16
	Routing-Instance	1/2636	521	string Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	16

Table 36: Logging Dictionary for Template Types (continued)

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
IPFlow TCP Timestamp	Retransmitted TCP packets uplink	1/2636	115	unsigned32	4
	Retransmitted TCP packets downlink	1/2636	116	unsigned32	4
	Smooth RTT uplink	1/2636	117	dateTimeMilliseconds	8
	Smooth RTT downlink	1/2636	118	dateTimeMilliseconds	8
	Client setup Time	1/2636	119	dateTimeMilliseconds	8
	Server Setup time	1/2636	120	dateTimeMilliseconds	8
	TCP flow creation timestamp	1/2636	121	dateTimeMilliseconds	8
	First Client Payload TS	1/2636	108	dateTimeMilliseconds	8
	Upload time	1/2636	113	dateTimeMilliseconds	8
	First Server Payload TS	1/2636	110	dateTimeMilliseconds	8
	Download time	1/2636	114	dateTimeMilliseconds	8
	Acknowledged volumes uplink	1/2636	122	unsigned64	8
		1/2636	123	unsigned64	8

Table 36: Logging Dictionary for Template Types (continued)

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
	Acknowledged volumes downlink				
IPFlow Timestamp	Flow Start Time	1/2636	101	dateTimeMilliseconds	8
	Flow End Time	1/2636	102	dateTimeMilliseconds	8
IPv4	Source IPv4 Address	0	8	ipv4Address	4
	Destination IPv4 Address	0	12	ipv4Address	4
IPv4 Extended (Not available if Next Gen Services with the MX-SPC3 services card are enabled)	IPv4 TOS/Class of Service	0	5	unsigned8	1
	IPv4 Source Mask	0	9	unsigned8	1
	IPv4 Destination Mask	0	13	unsigned8	1
	IPv4 Next Hop	0	15	ipv4Address	4
IPv6	Source IPv6 Address	0	27	ipv6Address	16
	Destination IPv6 Address	0	28	ipv6Address	16

Table 36: Logging Dictionary for Template Types (continued)

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
IPv6 Extended (Not available if Next Gen Services are enabled on the MX-SPC3 services card)	IPv6 Source Mask	0	29	unsigned8	1
	IPv6 Destination Mask	0	30	unsigned8	1
	IPv6 Next hop	0	62	ipv6Address	16
	Traffic Class	1/2636	126	unsigned8	1
L7 Application	Application Protocol	1/2636	151	string	32
	Application Name	1/2636	170	string	32
	Host	1/2636	157	string	64
Mobile Subscriber (Not available if Next Gen Services are enabled on the MX-SPC3 services card)	IMSI	1/2636	504	string	16
	MSISDN	1/2636	505	string	16
	IMEI	1/2636	506	string	16
	RAT-type	1/2636	507	unsigned8	1
	ULI	1/2636	508	string	13
	RADIUS Called Station ID	1/2636	509	string	32
PCC	PCC rule name	1/2636	901	string Not applicable if Next Gen Services are enabled.	64

Table 36: Logging Dictionary for Template Types (continued)

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
Status Code Distribution	Status code 1	1/2636	171	unsigned16	2
	Status code 2	1/2636	172	unsigned16	2
	Status code 3	1/2636	173	unsigned16	2
	Status code 4	1/2636	174	unsigned16	2
	Status code 5	1/2636	175	unsigned16	2
	Num Instances 1	1/2636	176	unsigned16	2
	Num Instances 2	1/2636	177	unsigned16	2
	Num Instances 3	1/2636	178	unsigned16	2
	Num Instances 4	1/2636	179	unsigned16	2
	Num Instances 5	1/2636	180	unsigned16	2

Table 36: Logging Dictionary for Template Types (continued)

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
Subscriber Data	NAS_IP_ADDR	1/2636	519	ipv4Address Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	4
	Subscriber Type	1/2636	515	unsigned8 1 for IP-based subscriber, 2 for IFL-based subscriber	1
	Subscriber IP address	1/2636	516	ipv4Address	4
	Subscriber VRF	1/2636	517	unsigned32 Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	4
	NAS Port ID	1/2636	518	string Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	32
	Accounting-Session-Id	1/2636	514	string Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	32

Table 36: Logging Dictionary for Template Types (continued)

Template Type	Field Name	Enterprise Bit/ID	Information Element Identifier	Data Type	Data Length (bytes)
	Class	1/2636	522	String Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	32
	NAS Port Type	1/2636	523	unsigned32 Not applicable for BNG subscribers, hence this value is not be honored (is filled with zero).	4
Transport Layer	Source Transport Port	0	7	unsigned16	2
	Destination Transport Port	0	11	unsigned16	2
Video (Not available if Next Gen Services are enabled on the MX-SPC3 services card)	Bitrate	1/2636	851	unsigned32	2
	Duration	1/2636	852	unsigned32	4
Wireline Subscriber (Not available if Next Gen Services are enabled on the MX-SPC3 services card)	UserName	1/2636	513	string	32

Configuring Logging and Reporting for Subscriber Management

To configure logging and reporting for traffic belonging to a subscriber, you configure LRF rules, collectors, and templates in an LRF profile; assign that LRF profile to the service set that is configured for application-aware policy control, and assign each LRF rule to a PCC rule to activate it.

NOTE: Starting in Junos OS Release 19.3R1, LRF is also supported for Broadband Subscriber Management if Next Gen Services are enabled on the MX-SPC3 services card).

To configure logging and reporting:

1. Install the LRF service package on any MS-MPC PICs that perform logging and reporting or on the MX-SPC3 services card if you have enabled if Next Gen Services.

See [“Installing Services Packages for Subscriber Management Logging and Reporting” on page 459](#).

2. Configure an LRF profile to specify a set of logging and reporting parameters, which includes data templates, collectors, and LRF rules.

See [“Configuring an LRF Profile for Subscribers” on page 460](#).

3. Assign the LRF profile to the service set that is configured for application-aware policy control.

See [“Applying Logging and Reporting Configuration to a Subscriber Management Service Set” on page 467](#).

4. Configure activation of an LRF rule with a static PCC rule.

See [“Configuring the Activation of an LRF Rule by a PCC Rule” on page 468](#).

RELATED DOCUMENTATION

[Logging and Reporting Function for Subscribers | 441](#)

Installing Services Packages for Subscriber Management Logging and Reporting

You must install the LRF service package on any MS-MPC PICs that perform logging and reporting or on an MX-SPC3 services card if you have enabled Next Gen Services.

To install the LRF service package:

1. Specify the MS-MPC PIC or MX-SPC3.

```
[edit chassis]
user@host# edit fpc slot-number pic pic-number
```

2. Install the services packages.

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider ]
user@host# set package jservices-lrf
```

RELATED DOCUMENTATION

| [Configuring Logging and Reporting for Subscriber Management](#) | 459

Configuring an LRF Profile for Subscribers

NOTE: Starting in Junos OS Release 19.3R1, LRF profiles are also supported for Broadband Subscriber Management if Next Gen Services are enabled on the MX-SPC3 services card.

Configure an LRF profile to specify a set of logging and reporting parameters, which includes data templates, collectors, and LRF rules.

To configure an LRF profile:

1. [Configuring the LRF Profile Name](#) | 461
2. [Configuring Policy-Based Logging](#) | 461
3. [\(Optional\) Configuring HTTP Transaction Logging](#) | 461
4. [Configuring Collectors](#) | 462
5. [Configuring Templates](#) | 463
6. [Configuring Logging and Reporting Rules](#) | 465

Configuring the LRF Profile Name

An LRF profile is identified by a name, which you later specify in the service set for the subscribers.

- Configure a name for the LRF profile.

```
[edit services lrf]
user@host# set profile profile-name
```

For example:

```
[edit services lrf]
user@host# set profile lrf_profile1
```

Configuring Policy-Based Logging

Policy-based logging causes the LRF rules to be activated by PCC rules in a static PCEF profile.

- Configure policy-based logging in the LRF profile.

```
[edit services lrf profile profile-name]
user@host# set policy-based-logging
```

For example:

```
[edit services lrf profile lrf_profile1]
user@host# set policy-based-logging
```

(Optional) Configuring HTTP Transaction Logging

Configure HTTP transaction logging if you want the HTTP metadata generated and sent separately for each transaction of a data session. This option is only relevant if the template specified in an LRF rule includes **http** in the **template-type**.

- Configure HTTP transaction logging in the LRF profile.

```
[edit services lrf profile profile-name]
user@host# set http-log-multiple-transactions
```

For example:

```
[edit services lrf profile lrf_profile1]
user@host# set http-log-multiple-transactions
```

Configuring Collectors

Configure one or more collectors that you want to receive logging and reporting data when an LRF rule is activated. You can configure up to eight collectors for an LRF profile. For each collector:

1. Configure a name for the collector.

```
[edit services lrf profile profile-name]
user@host# set collector collector-name
```

For example:

```
[edit services lrf profile lrf_profile1]
user@host# set collector collector1
```

2. Specify the destination IP address of the collector.

```
[edit services lrf profile profile-name collector collector-name destination]
user@host# set address collector-address
```

For example:

```
[edit services lrf profile lrf_profile1 collector collector1 destination]
user@host# set address 192.0.2.5
```

3. Specify the destination port of the collector.

```
[edit services lrf profile profile-name collector collector-name destination]
user@host# set port collector-port-number
```

For example:

```
[edit services lrf profile lrf_profile1 collector collector1 destination]
user@host# set port 4739
```

4. Configure the source address to be used when exporting data to the collector.

```
[edit services lrf profile profile-name collector collector-name]
user@host# set source-address source-address
```

For example:

```
[edit services lrf profile lrf_profile1 collector collector1]
user@host# set source-address 10.1.1.1
```

Configuring Templates

Configure one or more templates, each of which specifies a set of data to be transmitted when an LRF rule is activated. You can configure up to 16 templates for an LRF profile. For each template:

1. Configure a name for the template.

```
[edit services lrf profile profile-name]
user@host# set template template-name
```

For example:

```
[edit services lrf profile lrf_profile1]
user@host# set template template1
```

2. Configure a format for the template. Only the IPFIX format is supported for this release.

```
[edit services lrf profile profile-name template template-name]
user@host# set format ipfix
```

For example:

```
[edit services lrf profile lrf_profile1 template template1]
user@host# set format ipfix
```

3. Configure the template types, which specify the data fields to include. You must configure at least one type, and you can configure multiple types.

```
[edit services lrf profile profile-name template template-name]
user@host# set template-type template-type
```

For example:

```
[edit services lrf profile lrf_profile1 template template1]
user@host# set template-type http ipv4
```

This example results in a template that includes fields from both the HTTP and IPv4 templates.

NOTE: If you have enabled Next Gen Services on the MX-SPC3 services card, then the DNS, IFL subscriber, IPv4 extended, IPv6 extended, mobile subscriber, video, and wireline subscriber templates are not supported.

4. If you used the **ipflow-tcp-ts** template type, identify it as an IBM template to avoid a commit warning.

```
[edit services lrf profile profile-name]
user@host# set vendor-support ibm
```

5. Configure the interval, in seconds, at which you want the template to be retransmitted to the collector. The interval can be from 10 through 600, and the default is 60.

```
[edit services lrf profile profile-name template template-name]
user@host# set template-tx-interval tx-time
```

For example:

```
[edit services lrf profile lrf_profile1 template template1]
user@host# set template-tx-interval 100
```

6. Configure the type of trigger that causes the generation of data records and transmission to the collector. You can specify the trigger type as either the closing of the data session (default) or a data volume limit. The data volume limit value is specified within an LRF rule.

```
[edit services lrf profile profile-name template template-name]
user@host# set trigger-type (session-close | volume)
```

For example:

```
[edit services lrf profile lrf_profile1 template template1]
user@host# set trigger-type volume
```

Configuring Logging and Reporting Rules

Configure one or more LRF rules, which control how data sessions are logged and reported. You can configure up to 32 LRF rules for an LRF profile. For each LRF rule:

1. Configure a name for the LRF rule.

```
[edit services lrf profile profile-name]
user@host# set rule lrf-rule-name
```

For example:

```
[edit services lrf profile lrf_profile1]
user@host# set rule rule1
```

You cannot use the same LRF rule name in multiple LRF profiles.

2. Specify the collector that you want to receive the data if this rule is matched.

```
[edit services lrf profile profile-name rule lrf-rule-name ]
user@host# set then report collector collector-name
```

For example:

```
[edit services lrf profile lrf_profile1 rule rule1]
user@host# set then report collector collector1
```

3. Specify the template that identifies the type of data to report if this rule is matched.

```
[edit services lrf profile profile-name rule lrf-rule-name]
user@host# set then report template template-name
```

For example:

```
[edit services lrf profile lrf_profile1 rule rule1]
user@host# set then report template template1
```

4. If you specified **volume** for the template's trigger type in Step 6 of ["Configuring Templates" on page 463](#), configure the data volume limit to be used for reporting by this rule.

```
[edit services lrf profile profile-name rule lrf-rule-name]
```

```
user@host# set then report volume-limit volume
```

The data volume, in megabytes, can be from 1 through 1024.

For example:

```
[edit services lrf profile lrf_profile1 rule rule1]
user@host# set then report volume-limit 4
```

5. If you specified **time** for the template's trigger type in Step 6 of “Configuring Templates” on page 463, configure the time limit to be used for reporting by this rule.

```
[edit services lrf profile profile-name rule lrf-rule-name]
user@host# set then report time-limit time-interval
```

The time limit, in seconds, can be from 60 through 1800. The default is 300.

For example:

```
[edit services lrf profile lrf_profile1 rule rule1]
user@host# set then report time-limit 360
```

RELATED DOCUMENTATION

[Logging and Reporting Function for Subscribers | 441](#)

[Applying Logging and Reporting Configuration to a Subscriber Management Service Set | 467](#)

[Configuring the Activation of an LRF Rule by a PCC Rule | 468](#)

[Configuring Custom Application Signatures | 435](#)

Applying Logging and Reporting Configuration to a Subscriber Management Service Set

NOTE: Starting in Junos OS Release 19.3R1, LRF profiles are also supported for Broadband Subscriber Management if Next Gen Services are enabled on the MX-SPC3 services card.

To use an LRF profile, you must assign it to the service set that is configured for application-aware policy control.

To assign an LRF profile to subscribers:

- Assign the LRF profile to the service set.

```
[edit services service-set service-set-name]  
user@host# set lrf-profile profile-name
```

RELATED DOCUMENTATION

[Logging and Reporting Function for Subscribers | 441](#)

[Configuring an LRF Profile for Subscribers | 460](#)

[Applying Services to Subscriber-Aware Traffic with a Service Set](#)

[Configuring Logging and Reporting for Subscriber Management | 459](#)

[Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control | 424](#)

Configuring the Activation of an LRF Rule by a PCC Rule

NOTE: Starting in Junos OS Release 19.3R1, LRF rules are also supported for Broadband Subscriber Management if Next Gen Services are enabled on the MX-SPC3 services card.

NOTE: If you are using Junos OS Subscriber Aware, you must be in maintenance mode to make a change to a PCC action profile. (See *Changing PCEF Profiles, PCC Rules, PCC Rulebases, Diameter Profiles, Flow Descriptions, and PCC Action Profiles*).

NOTE: If you are using Junos OS Broadband Subscriber Management, you cannot make a change to a PCC action profile that is being used by subscribers. To modify the PCC action profile, you must first log off the subscribers that are using the PCC action profile.

Before you configure activation of an LRF rule by a PCC rule, you must:

- Configure the LRF rule in an LRF profile.
- Configure policy-based logging in the LRF profile.
- Configure the PCC rule.

You use a PCC rule's matching conditions to activate an LRF rule, which controls how data sessions are logged and reported. You identify the LRF rule in the PCC rule's action profile.

You can configure a PCC rule to activate an LRF rule for Junos OS Subscriber Aware or for Junos OS Broadband Subscriber Management, but you use a different CLI hierarchy level for each product.

- If you are using Junos OS Subscriber Aware, configure PCC rules at the **[edit unified-edge pcef]** hierarchy level.
- If you are using Junos OS Broadband Subscriber Management, configure PCC rules at the **[edit services pcef]** hierarchy level.

To configure a PCC rule to activate an LRF rule:

1. Identify the PCC action profile that is used in the PCC rule.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef]
```



```
user@host# show pcc-rules rule-name
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef]
user@host# show pcc-rules rule-name
```

For example:

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef]
user@host# show pcc-rules all-traffic
```

```
from {
  flows {
    all;
  }
}
then {
  pcc-action-profile all-traffic-action;
}
```

For Junos OS Broadband Subscriber Management:

NOTE: The **from** statement is not applicable for Next Gen Services MX-SPC3 services card.

```
[edit services pcef]
user@host# show pcc-rules all-traffic
```

```
from {
  flows {
    all;
  }
}
then {
  pcc-action-profile all-traffic-action;
}
```

2. Assign the LRF rule to the PCC action profile.

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-action-profiles profile-name]  
user@host# set logging-rule lrf-rule-name
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-action-profiles profile-name]  
user@host# set logging-rule lrf-rule-name
```

For example:

For Junos OS Subscriber Aware:

```
[edit unified-edge pcef pcc-action-profiles all-traffic-action]  
user@host# set logging-rule rule1
```

For Junos OS Broadband Subscriber Management:

```
[edit services pcef pcc-action-profiles all-traffic-action]  
user@host# set logging-rule rule1
```

RELATED DOCUMENTATION

[Logging and Reporting Function for Subscribers | 441](#)

[Configuring an LRF Profile for Subscribers | 460](#)

[Configuring Policy and Charging Control Rules | 418](#)



Configuring HTTP Redirect Services

Configuring Captive Portal Content Delivery Services for Redirected Subscribers | **473**

Configuring Captive Portal Content Delivery Services for Redirected Subscribers

IN THIS CHAPTER

- [HTTP Redirect Service Overview | 473](#)
- [Remote HTTP Redirect Server Operation Flow | 480](#)
- [Local HTTP Redirect Server Operation Flow | 482](#)
- [Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)
- [Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)
- [Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)
- [Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)
- [Adding Subscriber Information to HTTP Redirect URLs | 529](#)
- [How to Automatically Remove the HTTP Redirect Service After the Initial Redirect | 532](#)
- [Example: Configuring HTTP Redirect Services Using a Next-Hop Method and Attaching It to a Static Interface | 534](#)

HTTP Redirect Service Overview

IN THIS SECTION

- [Services-Card-Based Captive Portal | 476](#)
- [Routing Engine-Based Captive Portal | 477](#)
- [Converged Service Provisioning for HTTP Redirect Services | 477](#)
- [Static Service Provisioning for HTTP Redirect Services | 478](#)

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 on an external device. The captive portal is often the initial page a subscriber sees after logging in to a subscriber session. The captive portal also receives and manages HTTP requests to unauthorized Web resources.

For example, the user might be redirected to a webpage that shows a company logo and network usage policy or to a page where the subscriber pays for services. The captive portal typically provides authentication and authorization services for redirected subscribers before granting access to protected servers outside of a walled garden.

A *walled garden*, also known as a *white-list*, defines a group of servers where access is provided to subscribers without reauthorization through a captive portal. These walled gardens enable you to increase revenue by marketing various services to your customers.

Typical walled garden links are:

- Vendor services, such as automobile rentals
- Hotel and motel loyalty or corporate program portals
- Room services
- Local attractions and weather

NOTE: This documentation uses the terms *HTTP redirect service* and *captive portal content delivery (CPCD) service* interchangeably.

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 `cpdd` 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 response with a preconfigured redirect URL. The response message includes an HTTP status code. Starting in Junos OS Release 17.3R1, the status code that is returned depends on the HTTP version used by the HTTP client that sent the GET request. When the version is higher than HTTP 1.0, the redirect server returns the 307 (Temporary Redirect) status code. When the version is HTTP 1.0, the 302 (Found) status code is returned. In releases earlier than 17.3R1, the redirect server returns the 302 status code regardless of HTTP version. Both codes inform the HTTP client to use the original URL, rather than the redirect URL, for subsequent GET requests.

When the response to the HTTP request is sent to the subscriber, the original URL is preserved by optionally appending it to the end of the configured redirect URL. The maximum length of the redirect URL, including the appended original URL, is 128 bytes. Starting in Junos Release 17.3R1, the maximum length of the redirect URL is increased to 1360 bytes and the redirect server can append additional information about the subscriber to the redirect URL. The maximum length applies regardless of whether subscriber information is appended to the URL. To append the subscriber information, you can specify certain subscriber attributes

in the VSAs returned in the RADIUS Accept-Access message in response to the subscriber login or in a RADIUS Change of Authorization (CoA) message. This applies for both Activate-Service (26-65) and Deactivate-Service (26-66) VSAs. The subscriber information is retrieved from the subscriber session database.

The control handler maintains a connection with the cpcdd process on the Routing Engine to learn configuration changes, such as the redirect URL and the rewrite IP destination and port. To achieve faster performance, the control handler maintains a cache of relevant configured entities, such as URLs, on a Modular Port Concentrator (MPC).

HTTP redirect services are supported for both IPv4 and IPv6. You can attach an HTTP redirect service or service set to either a static or dynamic interface. For dynamic subscriber management, you can attach HTTP services or service sets dynamically at subscriber login or by using a RADIUS change of authorization (CoA).

Starting in Junos OS Release 17.2R1, there are three methods to configure HTTP redirect services. Starting in Junos OS Release 19.3R2, HTTP redirect can also be configured on the MX-SPC3 services processing card if Next Gen Services are enabled. [Table 37 on page 475](#) lists the methods supported for HTTP redirect services and the Junos OS releases that support each method.

BEST PRACTICE: We recommend that you use Junos OS Release 15.1 and higher releases to implement HTTP redirect services.

Table 37: Supported HTTP Redirect Methods by Release

Method		Junos OS Releases Supported
MS-DPC-based		(Not supported for Next Gen Services on the MX-SPC3 services card)
	Static	Releases earlier than 15.1
	Converged	Not supported
MS-MPC-based		(Not supported for Next Gen Services on the MX-SPC3 services card.)
	Static	Starting in Junos OS Release 15.1
	Converged	Starting in Junos OS Release 17.2
MX-SPC3-based		

Table 37: Supported HTTP Redirect Methods by Release (*continued*)

Method		Junos OS Releases Supported
	Static	Starting in Junos OS Release 19.3R2 if Next Gen Services are enabled on the MX-SPC3 services card.
	Converged	Starting in Junos OS Release 19.3R2 if Next Gen Services are enabled on the MX-SPC3 services card.
Routing Engine-based		
	Static	All Junos OS releases
	Converged	Starting in Junos OS Release 16.1R4 and 17.2

For all methods, you configure the walled garden as a static firewall service filter.

Services-Card-Based Captive Portal

IN THIS SECTION

- [MS-MPC-Based Captive Portal | 476](#)
- [MX-SPC3 Services Card-Based Captive Portal | 477](#)
- [Walled Garden Configured as a Service Filter | 477](#)

MS-MPC-Based Captive Portal

Starting in Junos OS Release 15.1R4, the only line card and interface card combination that supports HTTP redirect services on MX Series routers is the Multiservices Modular Port Concentrator (MS-MPC) with a Multiservices Modular Interface Card (MS-MIC). This combination provides improved scaling and high performance. MS-MICs and MS-MPCs have enhanced memory (16 GB for MS-MIC, 32 GB per NPU of MS-MPC) and processing capabilities. The services interfaces on MS-MPCs and MS-MICs are identified in the configuration with an *ms-* prefix (for example, *ms-1/2/1*).

NOTE: Throughout this documentation, the term *MS-MPC-based* refers to MPCs with MS-MICs installed and to MS-MICs alone when they are installed in MX Series routers that do not accept line cards.

MX-SPC3 Services Card-Based Captive Portal

Starting in Junos OS Release 19.3R2, you can configure HTTP redirect services if Next Gen Services are enabled on the MX-SPC3 services card. The services interfaces on MX-SPC3s are identified in the configuration with a vms- prefix (for example, vms-1/2/1).

Walled Garden Configured as a Service Filter

Packet flow for a services-card-based captive portal differs depending on how you configure the walled garden. HTTP traffic destined to servers within the walled garden does not flow to the services card. However, any HTTP traffic destined outside of the walled garden flows to the services card.

- For subscriber requests contained within the first packet of data traffic, the system expects TCP proxy to generate a TCP SYN flag causing the data handler to perform a rule lookup and apply those rules to HTTP data flows.
 - For an HTTP rewrite condition—If the IP destination address is not provided in the policy, the control handler looks up the IP destination address.
 - For an HTTP redirect condition—TCP proxy is triggered to complete its three-way handshake.
- For HTTP request packets.
 - For an HTTP rewrite condition—The control handler uses the cached IP destination address and modifies the data packet.
 - For an HTTP redirect condition—The control handler sends an HTTP 302 or 307 response with a preconfigured redirect URL.

Routing Engine-Based Captive Portal

The Routing Engine-based captive portal supports a walled garden as a firewall service filter for both static and converged services. As soon as the HTTP traffic matches the rules defined in the firewall service filter, the HTTP traffic is sent to the Routing Engine. The services interfaces on the Routing Engine are identified with an si- prefix (for example, si-1/1/0). The si- interface handles all redirect and rewrite traffic and services for the Routing Engine. The si- interface must be operational with a status of up to enable and activate the captive portal content delivery (CPCD) service. After the CPCD service is enabled, any change in the operational state of the si- interface does not affect existing CPCD services.

Converged Service Provisioning for HTTP Redirect Services

Starting in Junos OS Release 17.2R1, converged service provisioning is supported for both Routing Engine-Based and MS-MPC/MS-MIC-based captive portals. Starting in Junos OS Release 19.3R2, converged service provisioning is also supported for MX-SPC3 services card-based captive portals if Next Gen Services are enabled on the MX-SPC3 services card. Converged service provisioning means you can configure service provisioning in a dynamic profile. You can specify user-defined variables for services that are populated by means of a RADIUS VSA or a Change of Authorization (CoA) message.

For example, you might want to have a different redirect URL for each subscriber. You can create a `redirect-url` variable in the dynamic profile, then configure a service rule to redirect the matching subscriber to `$redirect-url`. When RADIUS authenticates the user, the Activate-Service VSA (26-65) provides the URL specific to that user.

Static Service Provisioning for HTTP Redirect Services

Starting in Junos OS Release 17.4R1, static service provisioning is supported for both Routing Engine-Based and MS-MPC/MS-MIC-based captive portals. Starting in Junos OS Release 19.3R2, static service provisioning is also supported for MX-SPC3-based captive portals if Next Gen Services are enabled on the MX-SPC3 services card. Static service provisioning means you can configure service provisioning in a static profile. You can specify user-defined variables (for example, `http://portal.wifi.example.com/xx?wlanuseraddr=%subsc-ip%&nasaddr=%nas-ip%&acname=%ac-name%&url=%dest-url%&userlocation=%nas-port-id%&usermac=%mac-sa%&session-id=%sess-id%&username=%user-name%&wlanuseraddrv6=%subsc-ipv6%`) for services that are populated by means of a RADIUS VSA or a Change of Authorization (CoA) message.

In static CPCD, attributes in a redirect URL are not sent in the Juniper Networks VSAs, Activate-Service (26-65) and Deactivate-Service (26-66). You can configure it as shown in the following example:

```
captive-portal-content-delivery {
  rule redirect {
    match-direction input;
    term t1 {
      then {
        redirect url;
      }
    }
  }
}
```

The tokens in the url such as “subsc-ip”, “nas-ip”, “ac-name” must be specified between “%” symbol. The order of tokens does not matter.

Following is a list of token with their significance:

- `%subsc-ip%`—private IP address of the subscriber.
- `%nas-ip%`—BNG IP address.
- `%ac-name%`—It will be empty for the BNG.
- `%dest-url%`—The original request url.
- `%nas-port-id%`—Used for subscriber. This parameter must include interface name, pvlan and cvlan. The interface name could be physical or virtual interface name. For example, `ge0/0/0` or `ae0`. The pvlan and cvlan range is 14095

- %mac-sa%—WLAN client MAC address.
- %sess-id%—session-id of subscriber.
- %user-name%—username of a subscriber.
- %subsc-ipv6%—subscriber IPv6 address (only IANA address). If IANA address is not specified for the subscriber, this field will be empty.

Release History Table

Release	Description
19.3R2	Starting in Junos OS Release 19.3R2, HTTP redirect can also be configured on the MX-SPC3 services processing card if Next Gen Services are enabled.
19.3R2	Starting in Junos OS Release 19.3R2, you can configure HTTP redirect services if Next Gen Services are enabled on the MX-SPC3 services card.
19.3R2	Starting in Junos OS Release 19.3R2, converged service provisioning is also supported for MX-SPC3 services card-based captive portals if Next Gen Services are enabled on the MX-SPC3 services card.
19.3R2	Starting in Junos OS Release 19.3R2, static service provisioning is also supported for MX-SPC3-based captive portals if Next Gen Services are enabled on the MX-SPC3 services card.
17.3R1	Starting in Junos OS Release 17.3R1, the status code that is returned depends on the HTTP version used by the HTTP client that sent the GET request.
17.3R1	Starting in Junos Release 17.3R1, the maximum length of the redirect URL is increased to 1360 bytes and the redirect server can append additional information about the subscriber to the redirect URL.
17.2R1	Starting in Junos OS Release 17.2R1, there are three methods to configure HTTP redirect services.
17.2R1	Starting in Junos OS Release 17.2R1, converged service provisioning is supported for both Routing Engine-Based and MS-MPC/MS-MIC-based captive portals.
17.2R1	Starting in Junos OS Release 17.4R1, static service provisioning is supported for both Routing Engine-Based and MS-MPC/MS-MIC-based captive portals.
15.1R4	Starting in Junos OS Release 15.1R4, the only line card and interface card combination that supports HTTP redirect services on MX Series routers is the Multiservices Modular Port Concentrator (MS-MPC) with a Multiservices Modular Interface Card (MS-MIC).

RELATED DOCUMENTATION

[Local HTTP Redirect Server Operation Flow | 482](#)

[Remote HTTP Redirect Server Operation Flow | 480](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

[Adding Subscriber Information to HTTP Redirect URLs | 529](#)

[How to Automatically Remove the HTTP Redirect Service After the Initial Redirect | 532](#)

Remote HTTP Redirect Server Operation Flow

You can use the remote HTTP redirect feature in configurations where the redirect server resides outside of the MX Series router and on a policy server, such as Session and Resource Control (SRC).

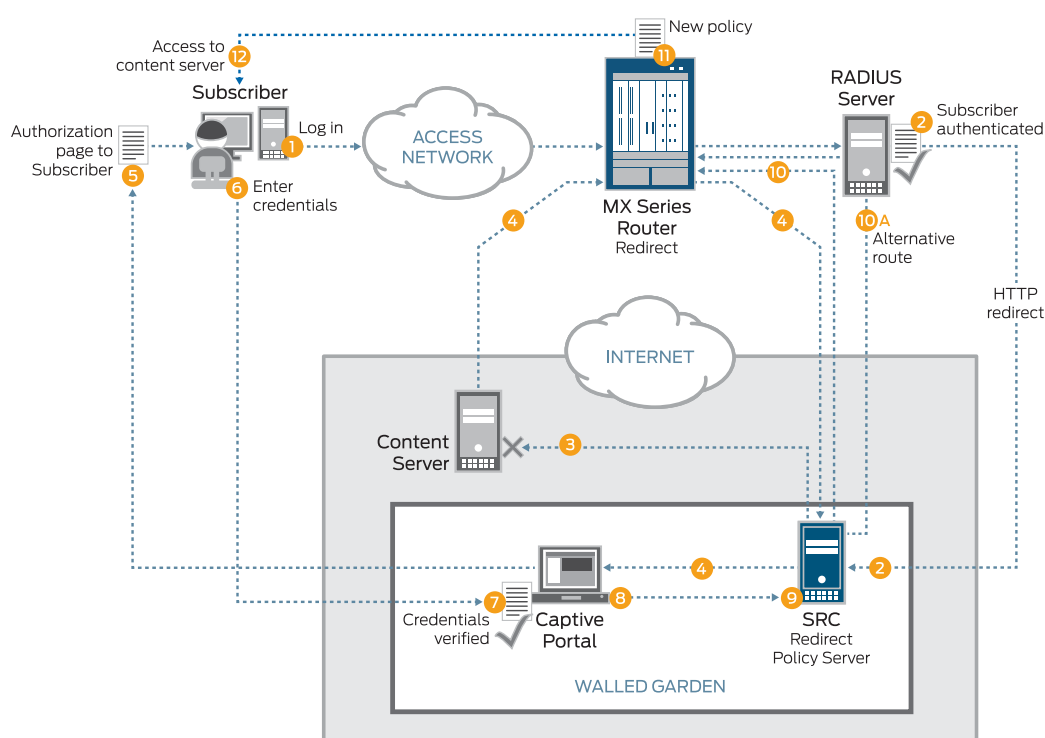
An HTTP redirect remote server that resides in a walled garden behind the router processes HTTP requests redirected to it and responds with a redirect URL that points to a captive portal. When you use a remote HTTP redirect server, you need to configure an HTTP service rule that rewrites the IP destination address of the incoming HTTP requests on the service router. The rewritten address ensures that the requests reach the remote HTTP redirect server before being redirected to a captive portal.

HTTP traffic is intercepted at the broadband network gateway (BNG) and the IP destination address is rewritten so that the HTTP requests are sent to the HTTP redirect server instead of the original destination. The HTTP redirect server sends a response with the HTTP 302 or 307 status code with the URL of the designated captive portal using either IPv4 destination address/destination port rewrite, or IPv6 destination address/destination port rewrite.

Figure 8 on page 481 shows the general service deployment during access configuration for a remote HTTP redirect server. The HTTP redirect server resides outside of the MX Series router on a policy server, such as SRC. Service attachment occurs at subscriber login, and service detachment occurs at subscriber logout.

NOTE: A complete HTTP redirect solution depends on back-end servers, such as SRC, captive portal, and RADIUS, and their integration specific to each customer's favored integration scheme.

Figure 8: Remote HTTP Redirect Server Deployment



1. The subscriber logs in connecting through the access network.
2. RADIUS authenticates the subscriber and sends a service activate (IP destination address rewrite), which redirects HTTP traffic to the redirect policy server (such as SRC) in a walled garden.
3. The subscriber attempts to access the content server (HTTP traffic).
4. The subscriber's HTTP traffic is first redirected to the SRC redirect policy server, which then 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's credentials.
8. The captive portal authorizes the subscriber and notifies the SRC redirect policy server.
9. The SRC redirect policy server checks the subscriber database and formulates a policy so the subscriber can access the content server.
10. The SRC redirect policy server sends the policy directly to the MX Series router using JSRC or Diameter.
Alternatively, the SRC redirect policy server notifies the RADIUS server, which in turn sends a change of authorization (CoA) to the MX Series router.
11. The MX Series router attaches the new policy, overriding the initial redirect policy.
12. The subscriber now gains access to the content server.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview | 473](#)

[Local HTTP Redirect Server Operation Flow | 482](#)

Local HTTP Redirect Server Operation Flow

You can use the local HTTP redirect feature in configurations where the redirect server resides locally on the MX Series router.

An HTTP redirect local server that resides locally on an MX Series router processes HTTP requests redirected to it and responds with a redirect URL that points 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.


A walled garden is a group of servers that provide subscriber access to sites within the walled garden without requiring reauthorization through a captive portal. HTTP request traffic from subscribers destined to servers outside of the walled garden is intercepted and redirected by either the captive portal content

delivery (CPCD) service or the Routing Engine. The CPCD service or Routing Engine locates the provisioned redirect URL for the specific subscriber and sends a response with the HTTP 302 or 307 status code that includes the located URL.

Figure 9 on page 483 shows the general service deployment during access configuration for a local HTTP redirect server. The HTTP redirect server resides locally on the MX Series router. Service attachment occurs at subscriber login, and service detachment occurs at subscriber logout.

NOTE: A complete HTTP redirect solution depends on back-end servers, such as SRC, captive portal, and RADIUS; their integration is specific to each customer's favored integration scheme.

Figure 9: Local HTTP Redirect Server Deployment



1. The subscriber logs in connecting through the access network.
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) outside the walled garden.

4. The subscriber's HTTP traffic is redirected to the captive portal by the MX Series 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 and notifies the SRC redirect policy server.
9. The SRC redirect policy server checks the subscriber database and formulates a policy so the subscriber can access the content server.
10. The SRC redirect policy server sends the policy directly to the MX Series router using JSRC or Diameter.
Alternatively, the SRC redirect policy server notifies the RADIUS server, which in turn sends a change of authorization (CoA) to the MX Series router.
11. The MX Series router attaches the new policy, overriding the initial redirect policy.
12. The subscriber now gains access to the content server.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview | 473](#)

[Remote HTTP Redirect Server Operation Flow | 480](#)

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

IN THIS SECTION

- [Configuring a Walled Garden as a Firewall Service Filter | 485](#)
- [Configuring HTTP Redirect for Local and Remote Redirect Servers | 488](#)
- [Configuring the Service Profile and the Service Set to Associate the Service Profile with a Service Interface | 490](#)

- [Attaching a CPCD Service Set and Service Filter to a Logical Interface | 491](#)
- [Installing a Service Package for CPCD Service | 492](#)

NOTE: Starting in Junos OS Release 19.3R2, static HTTP redirect service provisioning is also supported for MX-SPC3 services card-based captive portals if you have enabled Next Gen Services on the MX Series router.

A walled garden is a group of servers that provide subscriber access to sites within the walled garden without requiring reauthorization through a captive portal. The captive portal page is typically the initial page a subscriber sees after logging in to a subscriber session.

When subscribers try to access sites outside the walled garden, HTTP redirect services process IPv4 and IPv6 HTTP requests to manage that traffic. The subscriber HTTP request traffic that is not destined for the walled garden is sent to the redirect server, which responds with a redirect URL that sends traffic to a captive portal instead of to the unauthorized external site. The captive portal provides authentication and authorization services for the redirected subscribers before granting them access to protected servers outside of the walled garden.

The redirect server can be local or remote:

- Local redirect server—Resides on the router and redirects subscriber traffic to a captive portal inside a walled garden.
- Remote redirect server—Resides on a device such as a policy server inside a walled garden behind the router. The destination address for the subscriber's HTTP traffic is rewritten to the address of the remote redirect server. The remote server redirects subscriber traffic to a captive portal inside that walled garden.

You configure the walled garden as a firewall service filter. The service filter is attached to a static interface. The CPCD service is applied to a service interface (ms- on the MS-MPC or vms- on the MX-SPC3 services card) by means of a service set; the service set is then attached to a static interface.

Configuring a Walled Garden as a Firewall Service Filter

When you configure the walled garden as a firewall service filter, traffic that is destined to servers within the walled garden is identified and skipped. Because this traffic does not flow to the line card, handling requirements are reduced.

All other HTTP traffic is destined for addresses outside the walled garden. Because this traffic does not match the filter conditions, it flows to the line card for handling.

You can configure the service filter so that the walled garden contains a single server as the captive portal or a list of servers.

- Configure the walled garden with a single server as the captive portal:

1. Create the service filter.

```
[edit]
user@host# edit firewall family address-family service-filter filter-name
```

2. Define a filter term to identify and skip processing for traffic to the captive portal.

- a. Specify filter conditions to match traffic that is destined for the captive portal by specifying the destination address of the captive portal and the destination port.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-address ip-address
user@host# set term name from destination-port port-number
```

- b. Specify that the matching traffic skips processing on the line card.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

3. Define a filter term to identify HTTP traffic from all the traffic that did not match the previous term and send it for processing by CPCD service rules.

- a. Specify one or more HTTP port numbers to match the skipped HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-port http-port-number
```

- b. Specify that the matching traffic is processed by a CPCD service.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then service
```

4. Define a filter term to skip further action for any remaining, non-HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

For example, the following configuration creates a filter for IPv4 HTTP traffic, `walled-v4`, with the captive portal on 192.0.2.0. Traffic matching the address is skipped. Nonmatching traffic goes to term `http`, where HTTP traffic is picked out of all skipped traffic and sent to be processed according to a CPCD service. Finally, term `skip` causes all the remaining, non-HTTP traffic to be skipped.

```
[edit]
user@host# edit firewall family inet service-filter walled-v4
[edit firewall family inet service-filter walled-v4]
user@host# set term portal from destination-address 192.0.2.0
user@host# set term portal from destination-port 80
user@host# set term portal then skip
user@host# set term http from destination-port 80
user@host# set term http then service
user@host# set term skip then skip
```

- Configure the walled garden as a list or subnet of servers.

1. Create the service filter.

```
[edit]
user@host# edit firewall family address-family service-filter filter-name
```

2. Define a filter term.

- a. Specify filter conditions to match traffic that is destined for any server in the walled garden by specifying a destination prefix list of servers.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-prefix-list list-name
user@host# set term name from destination-port port-number
```

- b. Specify that the matching traffic skips processing on the line card.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

3. Define a filter term to identify HTTP traffic from all the traffic that did not match the previous term and send it for processing by CPCD service rules.

- a. Specify one or more HTTP port numbers to match the skipped HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-port http-port-number
```

- b. Specify that the matching traffic is processed by a CPCD service.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then service
```

4. Define a filter term to skip further action for any remaining, non-HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

5. (Optional) Define a prefix list that specifies servers within the walled garden. You can specify a subnet or multiple individual addresses.

```
[edit policy-options]
user@host# set prefix-list list-name ip-address/mask
user@host# set prefix-list list-name ip-address1
user@host# set prefix-list list-name ip-address2
```

For example, the following configuration creates a filter for IPv6 HTTP traffic, `walled-v6-list`, with a prefix list, `wg-list`, that specifies two servers in the walled garden. Filter term `portal6` identifies IPv6 traffic that is destined for the walled garden. Nonmatching traffic goes to term `http6`, where HTTP traffic is picked out of all skipped traffic and sent to be processed according to a CPCD service. Finally, term `skip` causes all the remaining, non-HTTP traffic to be skipped.

```
[edit]
user@host# edit firewall family inet6 service-filter walled-v6-list
user@host# set term portal6 from destination-prefix-list wg-list
user@host# set term portal6 then skip
user@host# set term http6 from destination-port [80 8080]
user@host# set term http6 then service
user@host# set term skip6 then skip
[edit policy-options]
user@host# set prefix-list wg-list 2001:db8::10.10
user@host# set prefix-list wg-list 2001:db8::10.22
```

Configuring HTTP Redirect for Local and Remote Redirect Servers

When HTTP requests are made for sites outside the walled garden, CPCD can redirect the traffic to a captive portal for authentication and authorization.

Configure a CPCD service rule that specifies the action to be taken for traffic destined outside the walled garden. This traffic was identified by the walled garden service filter and passed to the service. The action you configure depends on whether you are using a local or a remote HTTP redirect server:

- If you are using a local HTTP redirect server on the router, you specify the redirect action.
- If you are using a remote HTTP redirect server, which resides in a walled garden behind the router, then you cannot simply specify a redirect URL. In this case, the service rule must rewrite the IP destination address for the traffic. The new destination address is the address of the remote HTTP redirect server. The remote server then supplies a redirect URL to send the traffic to a captive portal.

The CPCD service is associated with a service interface by a service set. Both the service set and the walled garden service filter are applied to a statically configured interface.

1. Access the CPCD service configuration level.

```
[edit services]
user@host# edit captive-portal-content-delivery
```

2. Create a rule to apply to traffic destined outside the walled garden.

```
[edit services captive-portal-content-delivery]
edit rule name
```

3. Specify that the rule applies to incoming traffic.

```
[edit services captive-portal-content-delivery rule name]
user@host# set match-direction input
```

4. Define a rule term for CPCD to apply an action to HTTP traffic. Because the walled garden is configured as a service filter, the traffic is already filtered to be HTTP traffic before being sent to the service.

- For a local HTTP redirect server, provide the redirect URL, which is the URL of the captive portal with the original URL (outside the walled garden) appended:

```
[edit services captive-portal-content-delivery rule name]
user@host# set term name then redirect redirect-url/url=%dest-url%
```

- For a remote HTTP redirect server, provide the destination address of the remote server:

```
[edit services captive-portal-content-delivery rule name]
user@host# set term name then rewrite destination-address remote-server-address
```

NOTE: If you want the service to apply to both redirect and rewrite traffic, you can either configure a single rule with multiple terms to manage both cases, or separate rules for each case.

For example, in the following configuration for a local server, the CPCD service rule `redir-svc` redirects traffic to a captive portal, <http://www.portal.example.com>. The original URL entered by the subscriber is appended to the redirect URL.

```
user@host# edit services captive-portal-content-delivery
user@host# edit rule redir-svc
user@host# set match-direction input
user@host# set term redir1 then redirect http://www.portal.example.com/url=%dest-url%
```

The following configuration for a remote server creates CPCD service rule `rewr-svc` that rewrites the original destination address to the address of the remote server, 192.0.2.230.

```
user@host# edit services captive-portal-content-delivery
user@host# edit rule rewr-svc
user@host# set match-direction input
user@host# set term rewr1 then rewrite destination-address 192.0.2.230
```

Configuring the Service Profile and the Service Set to Associate the Service Profile with a Service Interface

Service sets define one or more services to be performed by the MS-MPC/MS-MIC, or the MX-SPC3 services card if you have enabled NEXT Gen Services on the MX Series router. For HTTP redirect services, you define a CPCD service profile that includes CPCD rules. The service set applies the CPCD service profile to a specific service interface.

1. Create the service profile.

```
[edit services captive-portal-content-delivery]
user@host# edit profile name
```

2. Specify one or more CPCD rules for the service profile.

```
[edit services captive-portal-content-delivery profile name]
user@host# set cpd-rules rule-name
```

3. Create the service set.

```
[edit services]
user@host# edit service-set name
```

4. Specify the CPCD service profile.

```
[edit services service-set name]
user@host# set captive-portal-content-delivery-profile name
```

5. Specify the service interface.

```
[edit services service-set name]
user@host# set interface-service service-interface interface-name
```

For example, the following configuration creates CPCD service profile `redir-prof`, which references the CPCD rule `redir-svc`. Service set `ss2` associates the CPCD service profile `redir-prof` with the service interface `ms-5/0/0`.

```
[edit services captive-portal-content-delivery]
user@host# edit profile redir-prof
user@host# set cpcd-rules redir-svc
[edit services]
user@host# edit service-set sset2
user@host# set captive-portal-content-delivery-profile redir-prof
user@host# set interface-service-service-interface ms-5/0/0
```

Attaching a CPCD Service Set and Service Filter to a Logical Interface

To use the HTTP redirect services, you must attach the CPCD service set to a logical interface. If the walled garden is configured as a service filter, then you must attach it to the same interface as the service set. Traffic arriving on and leaving that interface is filtered by the service filter. Traffic identified for servicing is sent to the MS-MPC, or to the MX-SPC3 services card if you have enabled Next Gen Services on the MX Series router, and the CPCD profile is applied at the service interface.

1. Configure the logical interface.

```
[edit interfaces interface-name]
user@host# edit unit logical-unit-number
```

2. Configure the address family.

```
[edit interfaces interface-name unit logical-unit-number]
user@host# edit family family
```

3. Configure the interface address.

```
[edit interfaces interface-name unit logical-unit-number family family]
user@host# set address address
```

4. Attach the service set and service filter to the interface.

```
[edit interfaces interface-name unit logical-unit-number family family]
user@host# set service input service-set set-name service-filter filter-name
user@host# set service output service-set set-name service-filter filter-name
```

For example, the following configuration attaches service set sset2 and service filter walled-v4 to ge-2/0/1.0 for the IPv4 address family. It assigns an address to the logical interface. The service set and filter are both applied to the interface input and output.

```
user@host# edit interfaces ge-2/0/1 unit 0 family inet
user@host# set address 203.0.113.5
user@host# set service input service-set sset2 service-filter walled-v4
user@host# set service output service-set sset2 service-filter walled-v4
```

Installing a Service Package for CPCD Service

To use CPCD services on an MS-MPC/MS-MIC, or on an MX-SPC3 services card if you have enabled Next Gen Services on the MX Series router, you configure a service interface on the MS-MIC or MX-SPC3. You must install the required service package on each MS-MIC that has a service interface or on the MX-SPC3 services card.

1. Configure Junos OS to support a service package on a service interface on an MX Series 5G Universal Routing Platform with MS-MPCs or an MX-SPC3 services card.

```
[edit]
user@host# edit chassis fpc slot-number pic number adaptive-services service-package
```


- 2. Configure the CPCD service package to run on the PIC. When the **extension-provider** statement is first configured, the PIC reboots.

NOTE: Static MS-MPC-based or MX-SPC3 services card-based CPCD requires the CPCD service package (jservices-cpcd).

```
[edit chassis fpc slot-number pic number adaptive-services service-package]
user@host# set extension-provider package jservices-cpcd
```

- 3. (Optional) Enable PIC system logging to record or view system log messages on the PIC. You can specify one or more facilities, each at a configurable severity level.

```
[edit chassis fpc 1 pic 0 adaptive-services service-package]
user@host# set extension-provider syslog facility severity
```

For example, the following configuration loads the CPCD services package on the MS-MPC in chassis slot 1 and the MS-MIC in slot 0 of the MPC. System log messages are generated for any daemon and for local external applications at all severity levels.

```
user@host# edit chassis fpc 1 pic 0 adaptive-services service-package
[edit chassis fpc 1 pic 0 adaptive-services service-package]
user@host# set extension-provider package jservices-cpcd
[edit chassis fpc 1 pic 0 adaptive-services service-package]
user@host# set extension-provider syslog daemon any
user@host# set extension-provider syslog external any
```

Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R2, static HTTP redirect service provisioning is also supported for MX-SPC3 services card-based captive portals if you have enabled Next Gen Services on the MX Series router.

RELATED DOCUMENTATION

[Remote HTTP Redirect Server Operation Flow | 480](#)

[Local HTTP Redirect Server Operation Flow | 482](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Adding Subscriber Information to HTTP Redirect URLs | 529](#)

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

IN THIS SECTION

- [Configuring a Walled Garden as a Firewall Service Filter | 495](#)
- [Configuring HTTP Redirect for Local and Remote Redirect Servers | 498](#)
- [Configuring Parameterization for the Redirect URL | 500](#)
- [Configuring the Service Set to Associate the Service Profile with a Service Interface | 501](#)
- [Attaching a CPCD Service Set and Service Filter to a Dynamic Logical Interface | 502](#)
- [Installing a Service Package for CPCD Service | 504](#)

You can configure converged HTTP redirect services on MS-MPCs/MS-MICs. Starting in Junos OS Release 19.3R1, you can also configure converged HTTP redirect service provisioning on the MX-SPC3 services card if you have enabled Next Gen Services on the MX Series router.

Converged service provisioning separates service definition from service instantiation. After a service is defined, a service can be dynamically instantiated at subscriber login or by using a change of authorization (CoA) mid-session. Service instantiation uses only the name of the defined service, hiding all service details from system operators. Converged service provisioning supports service parameterization, which corresponds to dynamic variables within dynamic profiles.

For converged HTTP redirect services, this means that you define the service and service rules within a dynamic profile. The CPCD service rules are created dynamically based on the variables configured in the dynamic profile.

Optionally, you can choose to parameterize the redirect URL by including defining a **redirect-url** variable in the dynamic profile. The value of the variable is provided by a RADIUS VSA during subscriber bring-up or with a Change of Authorization (CoA) message. This enables you to customize the redirect URLs for each subscriber. You can define a default value for the URL that is used if no value is provided by RADIUS.

You configure the walled garden as a firewall service filter. It filters traffic so that only HTTP traffic destined outside the walled garden is passed to the dynamic service for processing. Just as for static HTTP redirect services, a service profile contains the service rules. You configure a service set outside the dynamic profile to associate the CPCD service profile with a specific ms service interface on an MS-MPC/MS-MIC or a vsp service interface on an MX-SPC3 services card. Within the dynamic profile, you apply the service set and the walled garden service filter to a dynamic interface.

Configuring a Walled Garden as a Firewall Service Filter

When you configure the walled garden as a firewall service filter, traffic that is destined to servers within the walled garden is identified and skipped. Because this traffic does not flow to the line card, handling requirements are reduced.

All other HTTP traffic is destined for addresses outside the walled garden. Because this traffic does not match the filter conditions, it flows to the line card for handling.

You can configure the service filter so that the walled garden contains a single server as the captive portal or a list of servers.

- Configure the walled garden with a single server as the captive portal:

1. Create the service filter.

```
[edit]
user@host# edit firewall family address-family service-filter filter-name
```

2. Define a filter term to identify and skip processing for traffic to the captive portal.
 - a. Specify filter conditions to match traffic that is destined for the captive portal by specifying the destination address of the captive portal and the destination port.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-address ip-address
user@host# set term name from destination-port port-number
```

- b. Specify that the matching traffic skips processing on the line card.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

3. Define a filter term to identify HTTP traffic from all the traffic that did not match the previous term and send it for processing by CPCD service rules.
 - a. Specify one or more HTTP port numbers to match the skipped HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-port http-port-number
```

- b. Specify that the matching traffic is processed by a CPCD service.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then service
```

4. Define a filter term to skip further action for any remaining, non-HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

For example, the following configuration creates a filter for IPv4 HTTP traffic, `walled-v4`, with the captive portal on 192.0.2.0. Traffic matching the address is skipped. Nonmatching traffic goes to term `http`, where HTTP traffic is picked out of all skipped traffic and sent to be processed according to a CPCD service. Finally, term `skip` causes all the remaining, non-HTTP traffic to be skipped.

```
[edit]
user@host# edit firewall family inet service-filter walled-v4
[edit firewall family inet service-filter walled-v4]
user@host# set term portal from destination-address 192.0.2.0
user@host# set term portal from destination-port 80
user@host# set term portal then skip
user@host# set term http from destination-port 80
user@host# set term http then service
user@host# set term skip then skip
```

- Configure the walled garden as a list or subnet of servers.

1. Create the service filter.

```
[edit]
user@host# edit firewall family address-family service-filter filter-name
```

2. Define a filter term.

- a. Specify filter conditions to match traffic that is destined for any server in the walled garden by specifying a destination prefix list of servers.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-prefix-list list-name
```

```
user@host# set term name from destination-port port-number
```

- b. Specify that the matching traffic skips processing on the line card.

```
[edit firewall family inet service-filter filter-name]  
user@host# set term name then skip
```

3. Define a filter term to identify HTTP traffic from all the traffic that did not match the previous term and send it for processing by CPCD service rules.

- a. Specify one or more HTTP port numbers to match the skipped HTTP traffic.

```
[edit firewall family inet service-filter filter-name]  
user@host# set term name from destination-port http-port-number
```

- b. Specify that the matching traffic is processed by a CPCD service.

```
[edit firewall family inet service-filter filter-name]  
user@host# set term name then service
```

4. Define a filter term to skip further action for any remaining, non-HTTP traffic.

```
[edit firewall family inet service-filter filter-name]  
user@host# set term name then skip
```

5. (Optional) Define a prefix list that specifies servers within the walled garden. You can specify a subnet or multiple individual addresses.

```
[edit policy-options]  
user@host# set prefix-list list-name ip-address/mask  
user@host# set prefix-list list-name ip-address1  
user@host# set prefix-list list-name ip-address2
```

For example, the following configuration creates a service filter for IPv6 HTTP traffic, `walled-v6-list`, with a prefix list, `wg-list`, that specifies two servers in the walled garden. Filter term `portal6` identifies IPv6 traffic that is destined for the walled garden. Nonmatching traffic goes to term `http6`, where HTTP traffic is picked out of all skipped traffic and sent to be processed according to a CPCD service. Finally, term `skip` causes all the remaining, non-HTTP traffic to be skipped.

```
[edit]  
user@host# edit firewall family inet6 service-filter walled-v6-list
```

```

user@host# set term portal6 from destination-prefix-list wg-list
user@host# set term portal6 then skip
user@host# set term http6 from destination-port [80 8080]
user@host# set term http6 then service
user@host# set term skip6 then skip
[edit policy-options]
user@host# set prefix-list wg-list 2001:db8::10.10
user@host# set prefix-list wg-list 2001:db8::10.22

```

Configuring HTTP Redirect for Local and Remote Redirect Servers

When HTTP requests are made for sites outside the walled garden, CPCD can redirect the traffic to a captive portal for authentication and authorization.

Configure a CPCD service rule that specifies the action to be taken for the HTTP traffic identified by the walled garden service filter and passed to the service. The action you configure depends on whether you are using a local or a remote HTTP redirect server:

- If you are using a local HTTP redirect server on the router, you specify the redirect action.
- If you are using a remote HTTP redirect server, which resides in a walled garden behind the router, then you cannot simply specify a redirect URL. In this case, the service rule must rewrite the IP destination address for the traffic. The new destination address is the address of the remote HTTP redirect server. The remote server then supplies a redirect URL to send the traffic to a captive portal.

1. Configure the dynamic profile.

```

[edit]
user@host# edit dynamic-profiles profile-name

```

2. Access the dynamic CPCD service configuration level.

```

[edit dynamic-profiles profile-name]
user@host# edit services captive-portal-content-delivery

```

3. Create a rule to apply to traffic destined outside the walled garden.

```

[edit dynamic-profiles profile-name services captive-portal-content-delivery]
edit rule name

```

4. Specify that the rule applies to incoming traffic.

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule name]
user@host# set match-direction input
```

5. Specify the action to take for the matching traffic .Because the walled garden is a service filter, the traffic is already identified as HTTP traffic before being sent to the service.

- For a local HTTP redirect server, provide the redirect URL, which is the URL of the captive portal with the original URL (outside the walled garden) appended:

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule name]
user@host# set term name then redirect redirect-url/url=%dest-url%
```

- For a remote HTTP redirect server, provide the destination address of the remote server:

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule name]
user@host# set term name then rewrite destination-address remote-server-address
```

NOTE: If you want the service to apply to both redirect and rewrite traffic, you can either configure a single rule with multiple terms to manage both cases, or separate rules for each case.

For example, in the following configuration for a local server, the dynamic profile `http-redir-converged` includes the CPCD service rule `redir-svc`. The rule redirects traffic to a captive portal, `http://www.portal.example.com`. The original URL entered by the subscriber is appended to the redirect URL. The CPCD service profile `redir-prof` includes the rule, and will later be applied to a service interface by a service set.

```
user@host# edit dynamic-profiles http-redir-converged
user@host# edit services captive-portal-content-delivery
user@host# edit rule redir-svc
user@host# set match-direction input
user@host# set term redir1 then redirect http://www.portal.example.com/url=%dest-url%
```

The following configuration for a remote server creates CPCD service rule `rewr-svc` that rewrites the original destination address to the address of the remote server, `192.0.2.230`.

```
user@host# edit dynamic-profiles http-redir-converged
user@host# edit services captive-portal-content-delivery
user@host# edit rule rewr-svc
```

```
user@host# set match-direction input
user@host# set term rewr1 then rewrite destination-address 192.0.2.230
```

Configuring Parameterization for the Redirect URL

You can optionally choose to parameterize the redirect URL and the rewrite destination address by specifying user-defined variables in the dynamic profile. Parameterizing means that URL or address becomes a dynamic variable. The value is provided by RADIUS when the subscriber is authenticated or when a CoA is received. Consequently, you can use the RADIUS attributes to provide different URLs or destination addresses for different subscribers.

1. Configure the dynamic profile.

```
[edit]
user@host# edit dynamic-profiles profile-name
```

2. Access the custom variable configuration level.

```
[edit dynamic-profiles profile-name]
user@host# edit variables
```

3. Define the variable for the redirect URL, the rewrite destination address, or both. Specify that the value for the dynamic variable is provided by an external server, typically RADIUS.

NOTE: You can name the variables anything you like, but names like `redirect-url` and `rewrite-da` make the purpose clear.

```
[edit dynamic-profiles profile-name variables]
set variable-name mandatory
```

4. In the CPCD rule, specify the variable by prepending a dollar sign (\$) to the variable name.

- For a local HTTP redirect server, provide the redirect variable:

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule name]
user@host# set term name then redirect $variable-name
```

- For a remote HTTP redirect server, provide the destination address variable:


```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule name]
user@host# set term name then rewrite $variable-name
```

For example, the following configuration shows two user-defined variables, `redirect-url` and `rewrite-da` that require externally provided values when they are instantiated. CPCD service rule `redir1` specifies traffic is redirected to `$redirect-url`. CPCD service rule `rewr1` specifies that the destination address for the traffic is rewritten to `$rewrite-da`.

```
user@host# edit dynamic-profiles http-redir-converged
user@host# edit variables
user@host# set redirect-url mandatory
user@host# set rewrite-da mandatory
user@host# edit services captive-portal-content-delivery
user@host# edit rule redir-svc
user@host# set match-direction input
user@host# set term redir1 then redirect $redirect-url
user@host# edit rule rewr-svc
user@host# set match-direction input
user@host# set term rewr1 then rewrite $rewrite-da
```

Configuring the Service Set to Associate the Service Profile with a Service Interface

Service sets define one or more services to be performed by the MS-MPC/MS-MIC, or by the MX-SPC3 services card if you have enabled Next Gen Services on the MX Series router. For HTTP redirect services, you define a CPCD service profile that includes CPCD rules. The service set applies the CPCD service profile to a specific service interface.

1. Create the service profile.

```
[edit services captive-portal-content-delivery]
user@host# edit profile name
```

2. Specify one or more CPCD rules configured in the CPCD dynamic profile for the service profile.

```
[edit services captive-portal-content-delivery profile name]
user@host# set cpcd-rules rule-name
```

3. Specify that this is a converged CPCD service.

```
[edit services captive-portal-content-delivery profile name]
```

```
user@host# set dynamic
```

4. Create the service set.

```
[edit services]
user@host# edit service-set name
```

5. Specify the CPCD service profile.

```
[edit services service-set name]
user@host# set captive-portal-content-delivery-profile name
```

6. Specify the service interface.

```
[edit services service-set name]
user@host# set interface-service service-interface interface-name
```

For example, the following configuration creates CPCD service profile `redir-prof`, which references the CPCD rule `redir-svc`. Service set `cvgd` associates the CPCD service profile `rewr-prof` with the service interface `ms-2/0/0`.

```
[edit services captive-portal-content-delivery]
user@host# edit profile redir-prof
user@host# set cpd-rules redir-svc
user@host# set dynamic
[edit services]
user@host# edit service-set cvgd
user@host# set captive-portal-content-delivery-profile redir-prof
user@host# set interface-service-service-interface ms-2/0/0
```

Attaching a CPCD Service Set and Service Filter to a Dynamic Logical Interface

To use the HTTP redirect services, you must attach the CPCD service set to a logical interface. Because the walled garden is configured as a service filter, you must attach it to the same interface as the service set. Traffic arriving on and leaving that interface is filtered by the service filter. Traffic identified for servicing is sent to the MS-MPC or MX-SPC3 service interface where the CPCD profile is applied.

NOTE: This procedure shows only elements of the dynamic profile configuration that are specific to the converged services configuration. The complete dynamic profile depends on your use case.

1. Configure the dynamic profile.

```
[edit]
user@host# edit dynamic-profiles profile-name
```

2. Configure the dynamic physical interface.

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

3. Configure the dynamic logical interface.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name]
user@host# edit unit $junos-underlying-interface-unit
```

4. Configure the address family.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-underlying-interface-unit]
user@host# edit family family
```

5. Attach the service set and service filter to the interface.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-underlying-interface-unit
  family family]
user@host# set service input service-set set-name service-filter filter-name
user@host# set service output service-set set-name service-filter filter-name
```

For example, the following configuration creates the dynamic profile `http-redir-converged`. It specifies predefined variables to create the dynamic physical and logical interfaces in the IPv4 address family. The profile attaches service set `cvgd` and service filter `walled-v4` to the dynamic logical interface when it is created at subscriber login. The service set and filter are both applied to the interface input and output.

```

user@host# edit dynamic-profiles http-redir-converged
user@host# edit interfaces $junos-interface-ifd-name
user@host# edit unit $junos-underlying-interface-unit
user@host# edit family inet
user@host# set service input service-set cvgd service-filter walled-v4
user@host# set service output service-set cvgd service-filter walled-v4

```

Installing a Service Package for CPCD Service

To use CPCD services on an MS-MPC/MS-MIC, or on an MX-SPC3 services card if you have enabled USF on the MX Series router, you configure a service interface on the MS-MIC or MX-SPC3. You must install the required service packages on each MS-MIC that has a service interface or on an MX-SPC3.

1. Configure Junos OS to support a service package on a service interface on an MX Series 5G Universal Routing Platform with MS-MPCs or on an MX-SPC3 services card for Next Gen Services.

```

[edit]
user@host# edit chassis fpc slot-number pic number adaptive-services service-package

```

2. Configure the required service packages to run on the PIC. When the **extension-provider**

NOTE: Converged services MS-MPC-based or MX-SPC3-based CPCD requires both the CPCD service package (jservices-cpcd) and the mobile subscriber service package (jservices-mss).

```

[edit chassis fpc slot-number pic number adaptive-services service-package]
user@host# set extension-provider package jservices-cpcd,jservices-mss

```

3. (Optional) Enable PIC system logging to record or view system log messages on the PIC. You can specify one or more facilities, each at a configurable severity level.

```

[edit chassis fpc 1 pic 0 adaptive-services service-package]
user@host# set extension-provider syslog facility severity

```

For example, the following configuration loads the CPCD services package and the mobile subscriber services package on the MS-MPC in chassis slot 1 and the MS-MIC in slot 0 of the MPC. System log messages are generated for any daemon and for local external applications at all severity levels.

```

user@host# edit chassis fpc 1 pic 0 adaptive-services service-package
[edit chassis fpc 1 pic 0 adaptive-services service-package]
user@host# set extension-provider package jservices-cpcd,jservices-mss
[edit chassis fpc 1 pic 0 adaptive-services service-package]
user@host# set extension-provider syslog daemon any
user@host# set extension-provider syslog external any

```

Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R1, you can also configure converged HTTP redirect service provisioning on the MX-SPC3 services card if you have enabled Next Gen Services on the MX Series router.

RELATED DOCUMENTATION

Dynamic Profiles Overview

Dynamic Variables Overview

Junos OS Predefined Variables

User-Defined Variables

[HTTP Redirect Service Overview | 473](#)

[Remote HTTP Redirect Server Operation Flow | 480](#)

[Local HTTP Redirect Server Operation Flow | 482](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)

[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

[Adding Subscriber Information to HTTP Redirect URLs | 529](#)

Configuring Routing Engine-Based, Static HTTP Redirect Services

IN THIS SECTION

- [Configuring a Walled Garden as a Firewall Service Filter | 507](#)
- [Configuring HTTP Redirect for Local and Remote Redirect Servers | 510](#)

- [Configuring the Service Profile and the Service Set to Associate the Service Profile with a Service Interface | 511](#)
- [Attaching a CPCD Service Set and Service Filter to a Logical Interface | 513](#)
- [Inserting GET Header Tags That the HTTP Server Can Use to Control Content Access | 514](#)

NOTE: Starting in Junos OS Release 19.3R2, the HTTP redirect service is also supported if you have enabled Next Gen Services on the MX Series.

You can configure HTTP redirect services on the Routing Engine as an alternative to using an MS-MPC/MS-MIC or MX-SPC3 services card. You configure the walled garden as a firewall service filter. A walled garden is a group of servers that provide subscriber access to sites within the walled garden without requiring reauthorization through a captive portal. The walled garden service filter identifies traffic destined for the walled garden and traffic destined outside the walled garden. Only HTTP traffic destined outside the walled garden is sent to the Routing Engine for processing by the HTTP redirect service. The CPCD service is associated with a service interface on the Routing Engine by means of a service set.

The service interfaces on the Routing Engine are identified with an si- prefix (for example, si-1/1/0). The si- interface processes all redirect and rewrite traffic and services for the Routing Engine. The si- interface must be operational with a status of up to enable and activate the captive portal content delivery (CPCD) service. After the CPCD service is enabled, any change in the operational state of the si- interface does not affect existing CPCD services.

The CPCD service sends the subscriber HTTP request traffic that is not destined for the walled garden to a redirect server, which responds with a redirect URL. The redirect URL sends traffic to a captive portal instead of to the unauthorized external site. The captive portal provides authentication and authorization services for the redirected subscribers before granting them access to protected servers outside of the walled garden.

The redirect server can be local or remote:

- **Local redirect server**—Resides on the router and redirects subscriber traffic to a captive portal inside a walled garden.
- **Remote redirect server**—Resides on a device such as a policy server inside a walled garden behind the router. The destination address for the subscriber's HTTP traffic is rewritten to the address of the remote redirect server. The remote server redirects subscriber traffic to a captive portal inside that walled garden.

Configuring a Walled Garden as a Firewall Service Filter

When you configure the walled garden as a firewall service filter, traffic that is destined to servers within the walled garden is identified and skipped. All other HTTP traffic is destined for addresses outside the walled garden. Because this traffic does not match the filter conditions, it flows to the Routing Engine for handling.

You can configure the service filter so that the walled garden contains a single server as the captive portal or a list of servers.

- Configure the walled garden with a single server as the captive portal:

1. Create the service filter.

```
[edit]
user@host# edit firewall family address-family service-filter filter-name
```

2. Define a filter term to identify and skip processing for traffic to the captive portal.

- a. Specify filter conditions to match traffic that is destined for the captive portal by specifying the destination address of the captive portal and the destination port.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-address ip-address
user@host# set term name from destination-port port-number
```

- b. Specify that the matching traffic skips processing on the line card.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

3. Define a filter term to identify HTTP traffic from all the traffic that did not match the previous term and send it for processing by CPCD service rules.

- a. Specify one or more HTTP port numbers to match the skipped HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-port http-port-number
```

- b. Specify that the matching traffic is processed by a CPCD service.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then service
```

4. Define a filter term to skip further action for any remaining, non-HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

For example, the following configuration creates a filter for IPv4 HTTP traffic, `walled-v4`, with the captive portal on 192.0.2.0. Processing is skipped for traffic matching the address; the traffic is sent to the captive portal. Nonmatching traffic goes to term `http`, where HTTP traffic is picked out of all skipped traffic and sent to be processed according to a CPD service. Finally, term `skip` causes all the remaining, non-HTTP traffic to be skipped.

```
[edit]
user@host# edit firewall family inet service-filter walled-v4
[edit firewall family inet service-filter walled-v4]
user@host# set term portal from destination-address 192.0.2.0
user@host# set term portal from destination-port 80
user@host# set term portal then skip
user@host# set term http from destination-port 80
user@host# set term http then service
user@host# set term skip then skip
```

- Configure the walled garden as a list or subnet of servers.

1. Create the service filter.

```
[edit]
user@host# edit firewall family address-family service-filter filter-name
```

2. Define a filter term.

- a. Specify filter conditions to match traffic that is destined for any server in the walled garden by specifying a destination prefix list of servers.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-prefix-list list-name
user@host# set term name from destination-port port-number
```

- b. Specify that the matching traffic skips processing on the line card.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```


3. Define a filter term to identify HTTP traffic from all the traffic that did not match the previous term and send it for processing by CPCD service rules.

- a. Specify one or more HTTP port numbers to match the skipped HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-port http-port-number
```

- b. Specify that the matching traffic is processed by a CPCD service.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then service
```

4. Define a filter term to skip further action for any remaining, non-HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

5. (Optional) Define a prefix list that specifies servers within the walled garden. You can specify a subnet or multiple individual addresses.

```
[edit policy-options]
user@host# set prefix-list list-name ip-address/mask
user@host# set prefix-list list-name ip-address1
user@host# set prefix-list list-name ip-address2
```

For example, the following configuration creates a filter for IPv6 HTTP traffic, `walled-v6-list`, with a prefix list, `wg-list`, that specifies two servers in the walled garden. Filter term `portal6` identifies IPv6 traffic that is destined for the walled garden. Nonmatching traffic goes to term `http6`, where HTTP traffic is picked out of all skipped traffic and sent to be processed according to a CPCD service. Finally, term `skip6` causes all the remaining, non-HTTP traffic to be skipped.

```
[edit]
user@host# edit firewall family inet6 service-filter walled-v6-list
user@host# set term portal6 from destination-prefix-list wg-list
user@host# set term portal6 then skip
user@host# set term http6 from destination-port [80 8080]
user@host# set term http6 then service
user@host# set term skip6 then skip
[edit policy-options]
user@host# set prefix-list wg-list 2001:db8::10.10
user@host# set prefix-list wg-list 2001:db8::10.22
```

Configuring HTTP Redirect for Local and Remote Redirect Servers

When HTTP requests are made for sites outside the walled garden, CPCD can redirect the traffic to a captive portal for authentication and authorization.

Configure a CPCD service rule that specifies the action to be taken for traffic destined outside the walled garden. This traffic was identified by the walled garden service filter and passed to the service, or identified and accepted by the walled garden service rule. The action you configure depends on whether you are using a local or a remote HTTP redirect server:

- If you are using a local HTTP redirect server on the router, you specify the redirect action.
- If you are using a remote HTTP redirect server, which resides in a walled garden behind the router, then you cannot simply specify a redirect URL. In this case, the service rule must rewrite the IP destination address for the traffic. The new destination address is the address of the remote HTTP redirect server. The remote server then supplies a redirect URL to send the traffic to a captive portal.

The CPCD service is associated with a service interface by a service set. Both the service set and the walled garden service filter are applied to a statically configured interface.

1. Access the CPCD service configuration level.

```
[edit services]
user@host# edit captive-portal-content-delivery
```

2. Create a rule to apply to traffic destined outside the walled garden.

```
[edit services captive-portal-content-delivery]
edit rule name
```

3. Specify that the rule applies to incoming traffic.

```
[edit services captive-portal-content-delivery rule name]
user@host# set match-direction input
```

4. Define a rule term for CPCD to apply an action to HTTP traffic. Because the walled garden is configured as a service filter, the traffic is already filtered to be HTTP traffic before being sent to the service.

- For a local HTTP redirect server, provide the redirect URL, which is the URL of the captive portal with the original URL (outside the walled garden) appended:

```
[edit services captive-portal-content-delivery rule name]
user@host# set term name then redirect redirect-url/url=%dest-url%
```

- For a remote HTTP redirect server, provide the destination address of the remote server:

```
[edit services captive-portal-content-delivery rule name]
user@host# set term name then rewrite destination-address remote-server-address
```

NOTE: If you want the service to apply to both redirect and rewrite traffic, you can either configure a single rule with multiple terms to manage both cases, or separate rules for each case.

For example, in the following configuration for a local server, the CPCD service rule `redir-svc` redirects traffic to a captive portal, `http://www.portal.example.com`. The original URL entered by the subscriber is appended to the redirect URL.

```
user@host# edit services captive-portal-content-delivery
user@host# edit rule redir-svc
user@host# set match-direction input
user@host# set term redir1 then redirect http://www.portal.example.com/url=%dest-url%
```

The following configuration for a remote server creates CPCD service rule `rewr-svc` that rewrites the original destination address to the address of the remote server, `192.0.2.230`.

```
user@host# edit services captive-portal-content-delivery
user@host# edit rule rewr-svc
user@host# set match-direction input
user@host# set term rewr1 then rewrite destination-address 192.0.2.230
```

Configuring the Service Profile and the Service Set to Associate the Service Profile with a Service Interface

Service sets define one or more services to be performed by the Routing Engine. For HTTP redirect services, you define a CPCD service profile that includes CPCD rules. The service set applies the CPCD service profile to a specific service interface.

1. Create the service profile.

```
[edit services captive-portal-content-delivery]
user@host# edit profile name
```

2. Specify one or more CPCD rules for the service profile.

```
[edit services captive-portal-content-delivery profile name]
user@host# set cpcd-rules rule-name
```

3. Create the service set.

```
[edit services]
user@host# edit service-set name
```

4. Specify that the service set is for Routing Engine–Based CPCD.

```
[edit services service-set name]
user@host# set service-set-options routing-engine-services
```

5. Specify the CPCD service profile.

```
[edit services service-set name]
user@host# set captive-portal-content-delivery-profile name
```

6. Specify the service interface.

```
[edit services service-set name]
user@host# set interface-service service-interface interface-name
```

For example, the following configuration creates CPCD service profile `redir-prof`, which references the CPCD rule `redir-svc`. Service set `ss2` is specified as being for Routing-Engine-based CPCD. The set associates the CPCD service profile `redir-prof` with the service interface `si-4/0/0`.

```
[edit services captive-portal-content-delivery]
user@host# edit profile redir-prof
user@host# set cpcd-rules redir-svc
[edit services]
user@host# edit service-set ss2
user@host# set service-set-options routing-engine-service
user@host# set captive-portal-content-delivery-profile redir-prof
user@host# set interface-service-service-interface si-4/0/0
```

Attaching a CPCD Service Set and Service Filter to a Logical Interface

To use the HTTP redirect services, you must attach the CPCD service set to a logical interface. If the walled garden is configured as a service filter, then you must attach it to the same interface as the service set. Traffic arriving on and leaving that interface is filtered by the service filter. Traffic identified for servicing is sent to the Routing Engine service interface where the CPCD profile is applied.

1. Enable inline services and specify a bandwidth.

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

2. Configure the logical inline services interface.

```
[edit interfaces interface-name]
user@host# edit unit logical-unit-number
```

3. Configure the address family.

```
[edit interfaces interface-name unit logical-unit-number]
user@host# edit family family
```

4. Attach the service set and service filter to the interface.

- Static interface:

```
[edit interfaces interface-name unit logical-unit-number family family]
user@host# set service input service-set set-name service-filter filter-name
user@host# set service output service-set set-name service-filter filter-name
```

- Dynamic interface

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family family]
user@host# set service input service-set set-name service-filter filter-name
user@host# set service output service-set set-name service-filter filter-name
```

For example, the following configuration enables inline services on the line card in chassis slot 4 and on the MIC in slot 0 of the line card. It assigns an address to the logical interface. Then it attaches service set sset2 and service filter walled-v4 to ge-2/0/1.0 for the IPv4 address family. The service set and filter are both applied to the interface input and output.

```

user@host# edit chassis fpc 4 pic 0 inline-services bandwidth 1g
user@host# edit interfaces ge-2/0/1 unit 0 family inet
user@host# set address 203.0.113.5
user@host# set service input service-set sset2 service-filter walled-v4
user@host# set service output service-set sset2 service-filter walled-v4

```

Inserting GET Header Tags That the HTTP Server Can Use to Control Content Access

In some cases you might want your HTTP server to determine whether to allow users to access content. Starting in Junos OS Release 19.1, you can configure Routing Engine-based, static HTTP redirect service filters to specify tags that the Routing Engine inserts in to the packet header of HTTP GET messages for this purpose. You can insert tags for the router hostname or the subscriber's MAC address, IPv4 address, or IPv6 address.

The following steps correspond to [Figure 10 on page 515](#).

1. The user's device, the HTTP client, performs a TCP handshake sequence with the HTTP server.
2. When the handshake is successful, the client sends an HTTP GET with the URL requested by the user.
3. The Routing Engine modifies that URL by concatenating a string of random characters enclosed by /\$ and \$/. The string length matches the combined length of the tags that will be inserted later. The string serves as an identifier when returned by the client.


Suppose the length of the tags to be inserted is 30 characters, and the requested URL is `http://192.51.100.20/test.html`. The Routing Engine returns the URL modified with a string of 30 random characters as in the following example:

```
http://192.51.100.20/test.html/$IIGSbVdNDTDvnJFIAYoysXwVJawoYj$/$/
```

4. The Routing Engine sends the modified URL with a status code of 302 (Found) or 307 (Temporary Redirect). The code sent depends on the version of HTTP being used and the version of Junos OS on the BNG. Both codes indicate to the client that the access request needs to be resent with the modified URL.
5. The Routing Engine resets the TCP connection with the client and the server.
6. The client performs a TCP handshake with the HTTP server for a modified URL.
7. The client sends an HTTP GET with the modified URL.
8. The Routing Engine checks whether the length of the concatenated string is the same as it sent to the client.
 - If the length is correct, it strips the URL back to the original requested URL, inserts the tags into the GET header, and forwards the GET to the HTTP server. If configured, the GET can be optionally forwarded to a redirect URL instead of the original requested server.
 - If the length is not correct, the Routing Engine drops the packet and increments the drop counter.

9. The HTTP server evaluates the GET message and sends a response to the client with a status code of 200 (OK) if it grants access or 403 (Forbidden) if the request is rejected.
10. The Routing Engine terminates the TCP connection with the client and the server.

Figure 10: Tag Insertion for HTTP Redirect Message Flow.



NOTE: Tags are inserted into the header in the same order as they are configured. The tag name is case-sensitive so that **tag ABCD** and **tag abcd** are processed as different names.

To configure tags to be inserted in GET headers:

1. Access the CPCD service configuration level.

```
[edit services]
user@host# edit captive-portal-content-delivery
```

2. Create a rule to apply to traffic destined outside the walled garden.

```
[edit services captive-portal-content-delivery]
edit rule name
```

3. Specify that the rule applies to incoming traffic.

```
[edit services captive-portal-content-delivery rule name]
user@host# set match-direction input
```

4. (Optional) Specify one or more destination addresses to filter traffic for tagging.

```
[edit services captive-portal-content-delivery rule name]
user@host# set term name from destination-address address
```

NOTE: Alternatively, you can specify destination addresses for identifying traffic in the firewall service filter.

5. Define a rule term for CPCD to apply an action to HTTP traffic. Because the walled garden is configured as a service filter, the traffic is already filtered to be HTTP traffic before being sent to the service.

```
[edit services captive-portal-content-delivery rule name]
user@host# set term name then insert tag tag-name tag-value tag-value
user@host# set term name then insert tag tag-name tag-value tag-value
```

For example, the following configuration creates a service rule, insert-rule, that matches traffic on the input interface. Term t1 inserts two tags, x-mac-addr with the subscriber's MAC address and x-sub-ip with the value of the subscriber's IPv4 address.

```
[edit]
```



```

user@host# edit services captive-portal-content-delivery rule insert-rule
user@host# set match-direction input
user@host# set term t1 then insert tag x-mac-addr tag-value subscriber-mac-addr
user@host# set term t1 then insert tag x-sub-ip tag-value subscriber-ip

```

In the following sample rule, only traffic with a destination address that matches 198.51.100.50 or 198.51.100.75 is tagged. Tags are inserted for the subscriber's IP address and the hostname of the router. A second term in the rule provides a redirect URL where the traffic is forwarded instead of being sent to the original requested URL.

```

user@host# edit services captive-portal-content-delivery
user@host# set match-direction input
user@host# set rule tag-redirect term t1 from destination-address 198.51.100.50
user@host# set rule tag-redirect term t1 from destination-address 198.51.100.75
user@host# set rule tag-redirect term t1 then insert tag x-sub-ip tag-value subscriber-ip
user@host# set rule tag-redirect term t1 then insert tag x-hostname tag-value hostname
user@host# set rule tag-redirect term t2 then redirect http://www.portal.example.com
user@host# set profile http-insert-redirect cpdc-rules tag-redirect

```

As with any CPCD service rules for Routing-Engine-Based HTTP redirect, you must include the rules in a CPCD service profile, then use a CPCD service set to associate the profile with an inline service interface. The Routing Engine uses the rules to process HTTP traffic passed by a service filter on the same logical interface as the service set.

Consider the following sample configuration. The tag-redirect rule is defined to match traffic on the input interface and then insert two tags in the GET header, the value of the subscriber's IP address and the hostname of the router. The rule then provides a redirect URL for the tagged traffic. The CPCD service profile http-insert-redirect is defined to include this rule.

```

user@host# edit services captive-portal-content-delivery
user@host# set match-direction input
user@host# set rule tag-redirect term t1 then insert tag x-sub-ip tag-value subscriber-ip
user@host# set rule tag-redirect term t1 then insert tag x-hostname tag-value hostname
user@host# set rule tag-redirect term t2 then redirect http://www.portal.example.com
user@host# set profile http-insert-redirect cpdc-rules tag-redirect

```

The service set sset1 is defined as being for Routing Engine-based CPCD. It applies the CPCD service profile to an inline service interface.

```

user@host# edit services service-set sset1
user@host# set service-set-options routing-engine-services
user@host# set captive-portal-content-delivery-profile http-insert-redirect

```

```
user@host# set interface-service service-interface si-1/1/0
```

The service filter walled-tag identifies and acts on three kinds of traffic: HTTP traffic to send to the walled garden at 192.0.2.100, HTTP traffic destined for 198.51.100.50 to go to service processing, and all other traffic to be skipped. This is an example of matching a destination address in the service filter instead of in the service rule.

```
user@host# edit firewall family inet service-filter walled-tag
user@host# set term portal from destination-address 192.0.2.100
user@host# set term portal from destination-port 80
user@host# set term portal then skip
user@host# set term http-tag from destination-address 198.51.100.50
user@host# set term http-tag from destination-port 80
user@host# set term http-tag then service
user@host# set term skip then skip
```

The service-set sset1 and service filter walled-tag are applied to a logical interface.

```
user@host# edit chassis fpc 4 pic 0 inline-services bandwidth 1g
user@host# edit interfaces ge-2/0/1 unit 0 family inet
user@host# set address 203.0.113.5
user@host# set service input service-set sset1 service-filter walled-tag
user@host# set service output service-set sset1 service-filter walled-tag
```

Release History Table

Release	Description
19.1	Starting in Junos OS Release 19.1, you can configure Routing Engine-based, static HTTP redirect service filters to specify tags that the Routing Engine inserts in to the packet header of HTTP GET messages for this purpose.

RELATED DOCUMENTATION

- [HTTP Redirect Service Overview | 473](#)
- [Remote HTTP Redirect Server Operation Flow | 480](#)
- [Local HTTP Redirect Server Operation Flow | 482](#)
- [Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)
- [Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)
- [Adding Subscriber Information to HTTP Redirect URLs | 529](#)

Configuring Routing Engine-Based, Converged HTTP Redirect Services

IN THIS SECTION

- [Configuring a Walled Garden as a Firewall Service Filter | 520](#)
- [Configuring HTTP Redirect for Local and Remote Redirect Servers | 523](#)
- [Configuring Parameterization for the Redirect URL | 525](#)
- [Configuring the Service Set to Associate the Service Profile with a Service Interface | 526](#)
- [Attaching a CPCD Service Set and Service Filter to a Dynamic Logical Interface | 528](#)

NOTE: Starting in Junos OS Release 19.3R1, the HTTP redirect service is also supported if you have enabled Next Gen Services on the MX Series.

You can configure converged HTTP redirect services on the Routing Engine as an alternative to using an MS-MPC/MS-MIC or MX-SPC3 services card. Converged service provisioning separates service definition from service instantiation. After a service is defined, a service can be dynamically instantiated at subscriber login or by using a change of authorization (CoA) mid-session. Service instantiation uses only the name of the defined service, hiding all service details from system operators. Converged service provisioning supports service parameterization, which corresponds to dynamic variables within dynamic profiles.

For converged HTTP redirect services, this means that you define the service and service rules within a dynamic profile. The CPCD service rules are created dynamically based on the variables configured in the dynamic profile.

Optionally, you can choose to parameterize the redirect URL by including defining a **redirect-url** variable in the dynamic profile. The value of the variable is provided by a RADIUS VSA during subscriber bring-up or with a Change of Authorization (CoA) message. This enables you to customize the redirect URLs for each subscriber. You can define a default value for the URL that is used if no value is provided by RADIUS.

You configure the walled garden as a firewall service filter. A walled garden is a group of servers that provide subscriber access to sites within the walled garden without requiring reauthorization through a captive portal. The walled garden service filter identifies traffic destined for the walled garden and traffic destined outside the walled garden. Only HTTP traffic destined outside the walled garden is passed to the dynamic service for processing.

The service interfaces on the Routing Engine are identified with an si- prefix (for example, si-1/1/0). The si- interface processes all redirect and rewrite traffic and services for the Routing Engine. The si- interface

must be operational with a status of up to enable and activate the captive portal content delivery (CPCD) service. After the CPCD service is enabled, any change in the operational state of the si- interface does not affect existing CPCD services.

Just as for static HTTP redirect services, a service profile contains the service rules. You configure a service set outside the dynamic profile to associate the CPCD service profile with a specific si service interface on the Routing Engine. Within the dynamic profile, you apply the service set and the walled garden service filter to a dynamic interface.

Configuring a Walled Garden as a Firewall Service Filter

When you configure the walled garden as a firewall service filter, traffic that is destined to servers within the walled garden is identified and skipped. Because this traffic does not flow to the line card, handling requirements are reduced.

All other HTTP traffic is destined for addresses outside the walled garden. Because this traffic does not match the filter conditions, it flows to the line card for handling.

You can configure the service filter so that the walled garden contains a single server as the captive portal or a list of servers.

- Configure the walled garden with a single server as the captive portal:

1. Create the service filter.

```
[edit]
user@host# edit firewall family address-family service-filter filter-name
```

2. Define a filter term to identify and skip processing for traffic to the captive portal.
 - a. Specify filter conditions to match traffic that is destined for the captive portal by specifying the destination address of the captive portal and the destination port.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-address ip-address
user@host# set term name from destination-port port-number
```

- b. Specify that the matching traffic skips processing on the line card.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

3. Define a filter term to identify HTTP traffic from all the traffic that did not match the previous term and send it for processing by CPCD service rules.

- a. Specify one or more HTTP port numbers to match the skipped HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-port http-port-number
```

- b. Specify that the matching traffic is processed by a CPCD service.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then service
```

4. Define a filter term to skip further action for any remaining, non-HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

For example, the following configuration creates a filter for IPv4 HTTP traffic, `walled-v4`, with the captive portal on 192.0.2.0. Traffic matching the address is skipped. Nonmatching traffic goes to term `http`, where HTTP traffic is picked out of all skipped traffic and sent to be processed according to a CPCD service. Finally, term `skip` causes all the remaining, non-HTTP traffic to be skipped.

```
[edit]
user@host# edit firewall family inet service-filter walled-v4
[edit firewall family inet service-filter walled-v4]
user@host# set term portal from destination-address 192.0.2.0
user@host# set term portal from destination-port 80
user@host# set term portal then skip
user@host# set term http from destination-port 80
user@host# set term http then service
user@host# set term skip then skip
```

- Configure the `walled garden` as a list or subnet of servers.

1. Create the service filter.

```
[edit]
user@host# edit firewall family address-family service-filter filter-name
```

2. Define a filter term.

- a. Specify filter conditions to match traffic that is destined for any server in the `walled garden` by specifying a destination prefix list of servers.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-prefix-list list-name
user@host# set term name from destination-port port-number
```

- b. Specify that the matching traffic skips processing on the line card.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

3. Define a filter term to identify HTTP traffic from all the traffic that did not match the previous term and send it for processing by CPCD service rules.

- a. Specify one or more HTTP port numbers to match the skipped HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name from destination-port http-port-number
```

- b. Specify that the matching traffic is processed by a CPCD service.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then service
```

4. Define a filter term to skip further action for any remaining, non-HTTP traffic.

```
[edit firewall family inet service-filter filter-name]
user@host# set term name then skip
```

5. (Optional) Define a prefix list that specifies servers within the walled garden. You can specify a subnet or multiple individual addresses.

```
[edit policy-options]
user@host# set prefix-list list-name ip-address/mask
user@host# set prefix-list list-name ip-address1
user@host# set prefix-list list-name ip-address2
```

For example, the following configuration creates a service filter for IPv6 HTTP traffic, `walled-v6-list`, with a prefix list, `wg-list`, that specifies two servers in the walled garden. Filter term `portal6` identifies IPv6 traffic that is destined for the walled garden. Nonmatching traffic goes to term `http6`, where HTTP traffic is picked out of all skipped traffic and sent to be processed according to a CPCD service. Finally, term `skip` causes all the remaining, non-HTTP traffic to be skipped.

```
[edit]
user@host# edit firewall family inet6 service-filter walled-v6-list
user@host# set term portal6 from destination-prefix-list wg-list
user@host# set term portal6 then skip
user@host# set term http6 from destination-port [80 8080]
user@host# set term http6 then service
user@host# set term skip6 then skip
[edit policy-options]
user@host# set prefix-list wg-list 2001:db8::10.10
user@host# set prefix-list wg-list 2001:db8::10.22
```

Configuring HTTP Redirect for Local and Remote Redirect Servers

When HTTP requests are made for sites outside the walled garden, CPCD can redirect the traffic to a captive portal for authentication and authorization.

Configure a CPCD service rule that specifies the action to be taken for the HTTP traffic identified by the walled garden service filter and passed to the service. The action you configure depends on whether you are using a local or a remote HTTP redirect server:

- If you are using a local HTTP redirect server on the router, you specify the redirect action.
- If you are using a remote HTTP redirect server, which resides in a walled garden behind the router, then you cannot simply specify a redirect URL. In this case, the service rule must rewrite the IP destination address for the traffic. The new destination address is the address of the remote HTTP redirect server. The remote server then supplies a redirect URL to send the traffic to a captive portal.

1. Configure the dynamic profile.

```
[edit]
user@host# edit dynamic-profiles profile-name
```

2. Access the dynamic CPCD service configuration level.

```
[edit dynamic-profiles profile-name]
user@host# edit services captive-portal-content-delivery
```

3. Create a rule to apply to traffic destined outside the walled garden.

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery]
edit rule name
```

- Specify that the rule applies to incoming traffic.

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule name]
user@host# set match-direction input
```

- Specify the action to take for the matching traffic. Because the walled garden is a service filter, the traffic is already identified as HTTP traffic before being sent to the service.

- For a local HTTP redirect server, provide the redirect URL, which is the URL of the captive portal with the original URL (outside the walled garden) appended:

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule name]
user@host# set term name then redirect redirect-url/url=%dest-url%
```

- For a remote HTTP redirect server, provide the destination address of the remote server:

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule name]
user@host# set term name then rewrite destination-address remote-server-address
```

NOTE: If you want the service to apply to both redirect and rewrite traffic, you can either configure a single rule with multiple terms to manage both cases, or separate rules for each case.

For example, in the following configuration for a local server, the dynamic profile `http-redir-converged` includes the CPCD service rule `redir-svc`. The rule redirects traffic to a captive portal, `http://www.portal.example.com`. The original URL entered by the subscriber is appended to the redirect URL. The CPCD service profile `redir-prof` includes the rule, and will later be applied to a service interface by a service set.

```
user@host# edit dynamic-profiles http-redir-converged
user@host# edit services captive-portal-content-delivery
user@host# edit rule redir-svc
user@host# set match-direction input
user@host# set term redir1 then redirect http://www.portal.example.com/url=%dest-url%
```

The following configuration for a remote server creates CPCD service rule `rewr-svc` that rewrites the original destination address to the address of the remote server, `192.0.2.230`.

```
user@host# edit dynamic-profiles http-redir-converged
```



```

user@host# edit services captive-portal-content-delivery
user@host# edit rule rewr-svc
user@host# set match-direction input
user@host# set term rewr1 then rewrite destination-address 192.0.2.230

```

Configuring Parameterization for the Redirect URL

You can optionally choose to parameterize the redirect URL and the rewrite destination address by specifying user-defined variables in the dynamic profile. Parameterizing means that URL or address becomes a dynamic variable. The value is provided by RADIUS when the subscriber is authenticated or when a CoA is received. Consequently, you can use the RADIUS attributes to provide different URLs or destination addresses for different subscribers.

1. Configure the dynamic profile.

```

[edit]
user@host# edit dynamic-profiles profile-name

```

2. Access the custom variable configuration level.

```

[edit dynamic-profiles profile-name]
user@host# edit variables

```

3. Define the variable for the redirect URL, the rewrite destination address, or both. Specify that the value for the dynamic variable is provided by an external server, typically RADIUS.

NOTE: You can name the variables anything you like, but names like `redirect-url` and `rewrite-da` make the purpose clear.

```

[edit dynamic-profiles profile-name variables]
set variable-name mandatory

```

4. In the CPCD rule, specify the variable by prepending a dollar sign (\$) to the variable name.

- For a local HTTP redirect server, provide the redirect variable:

```

[edit dynamic-profiles profile-name services captive-portal-content-delivery rule name]

```

```
user@host# set term name then redirect $variable-name
```

- For a remote HTTP redirect server, provide the destination address variable:

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule name]
user@host# set term name then rewrite $variable-name
```

For example, the following configuration shows two user-defined variables, `redirect-url` and `rewrite-da` that require externally provided values when they are instantiated. CPCD service rule `redir1` specifies traffic is redirected to `$redirect-url`. CPCD service rule `rewr1` specifies that the destination address for the traffic is rewritten to `$rewrite-da`.

```
user@host# edit dynamic-profiles http-redir-converged
user@host# edit variables
user@host# set redirect-url mandatory
user@host# set rewrite-da mandatory
user@host# edit services captive-portal-content-delivery
user@host# edit rule redir-svc
user@host# set match-direction input
user@host# set term redir1 then redirect $redirect-url
user@host# edit rule rewr-svc
user@host# set match-direction input
user@host# set term rewr1 then rewrite $rewrite-da
```

Configuring the Service Set to Associate the Service Profile with a Service Interface

Service sets define one or more services to be performed by the Routing Engine. For HTTP redirect services, you define a CPCD service profile that includes CPCD rules. The service set applies the CPCD service profile to a specific service interface.

1. Create the service profile.

```
[edit services captive-portal-content-delivery]
user@host# edit profile name
```

2. Specify one or more CPCD rules configured in the CPCD dynamic profile for the service profile.

```
[edit services captive-portal-content-delivery profile name]
user@host# set cpcd-rules rule-name
```

3. Specify that this is a converged CPCD service.

```
[edit services captive-portal-content-delivery profile name]  
user@host# set dynamic
```

4. Create the service set.

```
[edit services]  
user@host# edit service-set name
```

5. Specify that the service set is for Routing Engine–Based CPCD.

```
[edit services service-set name]  
user@host# set service-set-options routing-engine-services
```

6. Specify the CPCD service profile.

```
[edit services service-set name]  
user@host# set captive-portal-content-delivery-profile name
```

7. Specify the service interface.

```
[edit services service-set name]  
user@host# set interface-service service-interface interface-name
```

For example, the following configuration creates CPCD service profile `redir-prof`, which references the CPCD rule `redir-svc`. Service set `cvgd` associates the CPCD service profile `rewr-prof` with the service interface `si-4/0/0`.

```
[edit services captive-portal-content-delivery]  
user@host# edit profile redir-prof  
user@host# set cpcd-rules redir-svc  
user@host# set dynamic  
[edit services]  
user@host# edit service-set cvgd  
user@host# set captive-portal-content-delivery-profile redir-prof  
user@host# set interface-service-service-interface si-4/0/0
```

Attaching a CPCD Service Set and Service Filter to a Dynamic Logical Interface

To use the HTTP redirect services, you must attach the CPCD service set to a logical interface. Because the walled garden is configured as a service filter, you must attach it to the same interface as the service set. Traffic arriving on and leaving that interface is filtered by the service filter. Traffic identified for servicing is sent to the Routing Engine service interface where the CPCD profile is applied.

NOTE: This procedure shows only elements of the dynamic profile configuration that are specific to the converged services configuration. The complete dynamic profile depends on your use case.

1. Configure the dynamic profile.

```
[edit]
user@host# edit dynamic-profiles profile-name
```

2. Configure the dynamic physical interface.

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

3. Configure the dynamic logical interface.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name]
user@host# edit unit $junos-underlying-interface-unit
```

4. Configure the address family.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-underlying-interface-unit]
user@host# edit family family
```

5. Attach the service set and service filter to the interface.

```
[edit dynamic-profiles profile-name interfaces $junos-interface-ifd-name unit $junos-underlying-interface-unit
  family family]
user@host# set service input service-set set-name service-filter filter-name
user@host# set service output service-set set-name service-filter filter-name
```

For example, the following configuration creates the dynamic profile `http-redirect-converged`. It specifies predefined variables to create the dynamic physical and logical interfaces in the IPv4 address family. The profile attaches service set `cvgd` and service filter `walled-v4` to the dynamic logical interface when it is created at subscriber login. The service set and filter are both applied to the interface input and output.

```
user@host# edit dynamic-profiles http-redirect-converged
user@host# edit interfaces $junos-interface-ifd-name
user@host# edit unit $junos-underlying-interface-unit
user@host# edit family inet
user@host# set service input service-set cvgd service-filter walled-v4
user@host# set service output service-set cvgd service-filter walled-v4
```

RELATED DOCUMENTATION

[Dynamic Profiles Overview](#)

[Dynamic Variables Overview](#)

[Junos OS Predefined Variables](#)

[User-Defined Variables](#)

[HTTP Redirect Service Overview | 473](#)

[Remote HTTP Redirect Server Operation Flow | 480](#)

[Local HTTP Redirect Server Operation Flow | 482](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Adding Subscriber Information to HTTP Redirect URLs | 529](#)

Adding Subscriber Information to HTTP Redirect URLs

NOTE: Starting in Junos OS Release 19.3R2, the HTTP redirect service is also supported if you have enabled Next Gen Services on the MX Series.

Starting in Junos OS Release 17.3R1, you can add subscriber information to a redirect URL to make it easier to track subscribers, change service policies, and provision services. For example, a WLAN service model might redirect subscribers to a captive portal when they connect to the network and open a browser. The captive portal may provide an opportunity to update or purchase new services or require subscribers

to enter their credentials before they can access a service. For example, the subscriber might be offered an opportunity to pay for a faster Internet connection.

You can configure the Juniper Networks RADIUS VSAs Activate-Service (26-65) or Deactivate-Service (26-66) to specify a format for the redirect URL that includes tokens for several subscriber attributes. The values for these tokens are retrieved from the subscriber session database and appended to the redirect URL. When the CPCD service is activated, the modified redirect URL is then returned to the requesting HTTP client in a message with an HTTP 302 or 307 status code. You can specify the tokens in any order. When the CPCD service is deactivated, the subscriber traffic is no longer redirected; the deactivation effectively removes the redirect rule for the subscriber,

When the subscriber subsequently logs in at the captive portal or purchases new services or updates, the web server hosting the captive portal confirms the action based on the supplied credentials. The server then contacts the RADIUS service to update the service policies for that particular subscriber. The subscriber attributes appended to the redirect URL enable RADIUS to determine exactly which subscriber to update. RADIUS then sends a CoA to the router to update the subscriber's policy and access.

[Table 38 on page 530](#) describes the supported subscriber tokens. If other tokens are included in the redirect URL format in the VSA, they are ignored.

Table 38: Supported subscriber Tokens for Redirect URLs

Token for URL Format	Subscriber Attribute
%subsc-ip%	Subscriber's private IP address.
%subsc-ipv6%	Subscriber's complete private IPv6 address (not just the prefix).
%nas-ip%	BNG IP address, configured with the router-id statement at the [edit routing-options] hierarchy level.
%ac-name%	This token is always empty on a BNG.
%dest-url%	Original, requested URL.
%nas-port-id%	Subscriber's interface information, contained in the RADIUS NAS-Port-Id attribute (87). The attribute must include the interface name (physical or logical) and the PVLAN or CVLAN identifiers. The VLAN identifiers are in the range 1 through 4095.
%mac-sa%	MAC address of the WLAN client (the device the subscriber uses to access the network).
%sess-id%	Subscriber session ID.

Table 38: Supported subscriber Tokens for Redirect URLs (*continued*)

Token for URL Format	Subscriber Attribute
%user-name%	Subscriber username.

NOTE: Refer to your RADIUS server documentation for information about configuring the service VSAs.

Configure the redirect URL with the desired tokens. In the following example, the redirect URL is `http://portal.wifi.example.com`. The tokens are delimited by the & (ampersand) character.

```
http://portal.wifi.example.com/xx?wlanuseraddr=%subsc-ip%
&nasaddr=%nas-ip%&url=%dest-url%&userlocation=%nas-port-id%
&usermac=%mac-sa%&acname=%ac-name%&session-id=%sess-id%
&username=%user-name%
```

The RADIUS service VSA includes the redirect URL with appended tokens in parentheses immediately following the name of the service to be activated—the dynamic service profile. In the following example, the profile is `http-redirect-converged2`:

```
http-redirect-converged2(http://portal.wifi.example.com/xx?wlanuseraddr=%subsc-ip%
&nasaddr=%nas-ip%&url=%dest-url%&userlocation=%nas-port-id%
&usermac=%mac-sa%&acname=%ac-name%&session-id=%sess-id%
&username=%user-name%
```

As an example, the returned redirect URL might look like the following when the tokens are replaced with the actual subscriber values retrieved from the session database:

```
http://portal.wifi.example.com?wlanuseraddr=192.0.2.66&nasaddr=203.0.113.1
&url=http%3A%2F%2F192.0.2.1%3A80%2Ftest.html&ip=192.0.2.1:80
&userlocation=ge-1/0/0:100&usermac=00:00:5E:00:53:42&acname=
&session-id=886&username=USER1@EXAMPLE.NET
```

You can configure adding subscriber information to the redirect URL for dynamic (converged) Routing Engine-based and dynamic MS-MPC/MS-MIC-based or MX-SPC3 services card-based CPCD.

Release History Table

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, you can add subscriber information to a redirect URL to make it easier to track subscribers, change service policies, and provision services.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview](#) | 473

How to Automatically Remove the HTTP Redirect Service After the Initial Redirect

In some deployments, you might want to always redirect your subscribers to the captive portal just once each session so that you can serve them advertisements or notifications. Thereafter, you want the subscribers to reach the URL that they specify without additional redirects.

In other deployments, HTTP redirect services might be set up so that the subscriber is redirected multiple times before being able to access the requested URL. For example, after logging in and requesting a URL, the subscriber is redirected to a payment page. After satisfying the payment requirements, the subscriber again requests the URL, but is redirected to an advertisement page, such as for more service offerings. The subscriber must request the URL again to reach the target. In this business case, you might want to simplify access for certain customers by removing the redirect service after the first redirect.

Removal of the redirect service typically requires action by the external policy server such as the PCRF or RADIUS server. For example, the RADIUS server might send a CoA message to deactivate the service. Starting in Junos OS Release 19.4R1, you can configure the router to automatically remove the redirect service when triggered. You can use this automatic method when you do not want the external policy server to be involved in removing the service. The trigger for automatic removal is the initial HTTP GET request from the subscriber. When the subscriber initially requests the URL, the subscriber is redirected once to the captive portal the first time the URL is requested. That Get request causes the router to remove the redirect service, so that the next request for the URL takes the subscriber directly to that location.

Use one of the following methods to configure the automatic removal feature:

- Enable automatic removal for static redirect services.

```
[edit services captive-portal-content-delivery]
user@host# set auto-deactivate initial-get
```

- Disable automatic removal for static redirect services.

```
[edit services captive-portal-content-delivery]
user@host# set auto-deactivate never
```

- Enable automatic removal for dynamic redirect services.

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery]
```



```
user@host# set auto-deactivate initial-get
```

- Disable automatic removal for dynamic redirect services.

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery]
user@host# set auto-deactivate never
```

For dynamic HTTP redirect services, you can also create a user-defined variable to enable or disable automatic removal. The variable value, **initial-get** or **never**, is supplied by either the external policy server or a default value that you define. To use the variable:

1. Specify the user-defined variable in the dynamic profile.

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery]
user@host# set auto-deactivate $variable-name
```

2. Configure your external policy server to provide the value. See your server documentation for information about how to do this.

3. (Optional) Define a default value for the variable.

```
[edit dynamic-profiles profile-name]
user@host# set variables variable-name default-value default-value
```

For example, the following configuration specifies that in the absence of information from the external server, the initial GET message triggers automatic removal of the redirect service.

```
[edit dynamic-profiles profile-name services]
user@host# set captive-portal-content-delivery auto-deactivate $remove-redirect-service
user@host# set variables remove-redirect-service default-value initial-get
```

Release History Table

Release	Description
19.4R1	Starting in Junos OS Release 19.4R1, you can configure the router to automatically remove the redirect service when triggered

RELATED DOCUMENTATION

[HTTP Redirect Service Overview | 473](#)[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

Example: Configuring HTTP Redirect Services Using a Next-Hop Method and Attaching It to a Static Interface

IN THIS SECTION

- [Requirements | 534](#)
- [Overview | 535](#)
- [Configuration | 535](#)
- [Verification | 549](#)

This example shows how to configure HTTP redirect services using a next-hop method and attaching it to a static interface.

Requirements

This example uses the following hardware and software components:

- MX240, MX480, or MX960 Universal Routing Platform with a Multiservices Modular PIC Concentrator (MS-MPC) and Multiservices Modular Interfaces Card (MS-MIC) installed.
- Junos OS Release 15.1 or later.

Before you begin:

- Configure the connection between the redirect server and the MX Series router.
- Define the source address (203.0.113.0/24 is used in this example).
- Define one or more interfaces used for subscriber traffic.

Overview

HTTP redirect and rewrite services are supported for both IPv4 and IPv6. You can attach an HTTP redirect service or service set to either a static or dynamic interface. For dynamic subscriber management, you can attach HTTP services or service sets dynamically at subscriber login or by using a change of authorization (CoA). Using a next-hop method, you can configure HTTP redirect services and attach it to a static interface.

Configuration

IN THIS SECTION

- [Configuring the CPCD Services and Attaching Service Set to Static Interface | 537](#)
- [Configuring the Package and Installation for CPCD | 539](#)
- [Configuring the Static Interface, HTTP Redirect Filters, and Interface Service Options | 540](#)
- [Configuring the Additional Routing Instance and Assigning Its Next-Hop Static Interfaces | 544](#)
- [Configuring the Interface-Specific Filters to Direct HTTP Traffic | 545](#)
- [Configuring the Policy Option and Statement to Use a Private Blocks Prefix List | 548](#)

To configure HTTP redirect services using a next-hop method and attach it to a static interface, perform these tasks:

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, and then copy and paste the commands into the CLI.

```
[edit]
edit services captive-portal-content-delivery
set rule redirect match-direction input
set rule redirect term REDIRECT then redirect http://redirection-portal/redirection/
set profile http-redirect cpdc-rules redirect
edit services service-set http-redirect-sset
set captive-portal-content-delivery-profile http-redirect
set next-hop-service inside-service-interface ms-11/1/0.1
set next-hop-service outside-service-interface ms-11/1/0.2
[edit]
edit chassis fpc 11 pic 1 adaptive-services service-package
set extension-provider package jservices-cpcd
set extension-provider syslog daemon none
```

```

set extension-provider syslog external none
set extension-provider syslog kernel none
set extension-provider syslog pfe none
[edit]
set interfaces ge-0/0/1 unit 900 description VLAN REDIRECT
set interfaces ge-0/0/1 unit 900 vlan-id 900
set interfaces ge-0/0/1 unit 900 family inet filter input FF_HTTP_REDIRECT_IN
set interfaces ge-0/0/1 unit 900 family inet address 203.0.113.250/30
edit interfaces ms-11/1/0 services-options open-timeout 4
edit interfaces ms-11/1/0 services-options close-timeout 2
edit interfaces ms-11/1/0 services-options inactivity-tcp-timeout 5
edit interfaces ms-11/1/0 services-options inactivity-non-tcp-timeout 5
edit interfaces ms-11/1/0 services-options session-timeout 5
edit interfaces ms-11/1/0 services-options tcp-tickles 0
set interfaces ms-11/1/0 unit 1 family inet
set interfaces ms-11/1/0 unit 1 service-domain inside
set interfaces ms-11/1/0 unit 2 filter output FF_CPCD_REDIRECT_OUTPUT
set interfaces ms-11/1/0 unit 2 family inet
set interfaces ms-11/1/0 unit 2 service-domain outside
[edit]
edit routing-instances CPCD_REDIRECT
set instance-type virtual-router
set interface ms-1/1/0.1
set interface ms-1/1/0.2
set routing-options static route 0.0.0.0/0 next-hop ms-1/1/0.1
set routing-options static route 203.0.113.0/24 next-hop ms-1/1/0.2
[edit]
edit firewall family inet
set filter FF_CPCD_REDIRECT_OUTPUT interface-specific
set filter FF_CPCD_REDIRECT_OUTPUT term One then count back-to-default
set filter FF_CPCD_REDIRECT_OUTPUT term One then routing-instance default
set filter FF_HTTP_REDIRECT_IN interface-specific
set filter FF_HTTP_REDIRECT_IN term ACCEPTED_PREFIXES from prefix-list User-PRIVATE-Blocks-01
set filter FF_HTTP_REDIRECT_IN term ACCEPTED_PREFIXES then next term
set filter FF_HTTP_REDIRECT_IN term HTTP from protocol tcp
set filter FF_HTTP_REDIRECT_IN term HTTP from destination-port http
set filter FF_HTTP_REDIRECT_IN term HTTP then count HTTP
set filter FF_HTTP_REDIRECT_IN term HTTP then forwarding-class best-effort
set filter FF_HTTP_REDIRECT_IN term HTTP then routing-instance CPCD_REDIRECT
[edit]
edit policy-options policy-statement User-PRIVATE-Blocks-01
set 203.0.113.0/24

```

Configuring the CPCD Services and Attaching Service Set to Static Interface

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. 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.

```
[edit services]
user@host# edit captive-portal-content-delivery
```

2. Configure the service filter as a walled garden by defining the rule the router references when applying this HTTP service.

```
[edit services captive-portal-content-delivery]
user@host# edit rule redirect
```

3. Specify that the rule matches traffic coming in on the interface.

```
[edit services captive-portal-content-delivery rule redirect]
user@host# match-direction input
```

4. Create the term match and action properties for the CPCD rule for the HTTP service.

```
[edit services captive-portal-content-delivery rule redirect]
user@host# set term REDIRECT then redirect http://redirection-portal/redirection/
```

5. Create the CPCD profile for the IP destination address to redirect the HTTP service.

```
[edit services captive-portal-content-delivery]
user@host# edit profile http-redirect
```

6. Specify the CPCD rule for the HTTP service.

```
[edit services captive-portal-content-delivery profile http-redirect]
user@host# set cpcd-rules redirect
```

7. Create the service set for the CPCD services.

```
[edit services service-set]
user@host# edit http-redirect-sset
```

8. Specify the CPCD profile for the service set.

```
[edit services service-set http-redirect-sset]
user@host# set captive-portal-content-delivery-profile http-redirect
```

9. Specify the interface name for the next-hop service for an inside and outside service interfaces and attach them to static interfaces.

```
[edit services service-set http-redirect-sset]
user@host# set next-hop-service inside-service-interface ms-11/1/0.1
user@host# set next-hop-service outside-service-interface ms-11/1/0.2
```

Results

From configuration mode, confirm your configuration by entering the **show services** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
root@host# show services
captive-portal-content-delivery {
  rule redirect {
    match-direction input;
    term REDIRECT {
      then {
        redirect http://redirection-portal/redirection/;
      }
    }
  }
  profile http-redirect {
    cpcd-rules redirect;
  }
}
service-set http-redirect-sset {
  captive-portal-content-delivery-profile http-redirect;
  next-hop-service {
    inside-service-interface ms-11/1/0.1;
    outside-service-interface ms-11/1/0.2;
  }
}
```

```
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Package and Installation for CPCD

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Configure Junos OS to support the service package on a service interface on an MX Series 5G Universal Routing Platform with MS-MPCs/MS-MICs.

```
[edit chassis]
user@host# edit fpc 11 pic 1 adaptive-services service-package
```

2. Configure the CPCD service package to run on the PIC. When the **extension-provider** statement is first configured, the PIC reboots.

```
[edit chassis fpc 11 pic 1 adaptive-services service-package]
user@host# set extension-provider package jservices-cpcd
```

3. Enable PIC system logging to record or view system log messages on the PIC but do not include daemon, external, kernel, or Packet Forwarding Engine processes.

```
[edit chassis fpc 11 pic 1 adaptive-services service-package extension-provider]
user@host# set extension-provider syslog daemon none
user@host# set extension-provider syslog external none
user@host# set extension-provider syslog kernel none
user@host# set extension-provider syslog pfe none
```

Results

From configuration mode, confirm your configuration by entering the **show chassis** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
root@host# show chassis
  fpc 11 {
    pic 1 {
```

```

adaptive-services {
  service-package {
    extension-provider {
      package jservices-cpcd;
      syslog {
        daemon none;
        external none;
        kernel none;
        pfe none;
      }
    }
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Static Interface, HTTP Redirect Filters, and Interface Service Options

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Configure a Gigabit interface with a logical interface on which traffic arrives before it is redirected.

```

[edit interfaces]
user@host# edit ge-0/0/1 unit 900

```

2. Assign a description and VLAN ID to the logical interface.

```

[edit interfaces ge-0/0/1 unit 900]
user@host# set description VLAN-REDIRECT
user@host# set vlan-id 900

```

3. Configure the IPv4 family for the interface.

```

[edit interfaces ge-0/0/1 unit 900]
user@host# edit family inet

```

4. Configure an input filter to evaluate when packets are received and redirected on the interface.


```
[edit interfaces ge-0/0/1 unit 900 family inet]
user@host# set filter input FF_HTTP_REDIRECT_IN
```

5. Configure an address for the input filter.

```
[edit interfaces ge-0/0/1 unit 900 family inet]
user@host# set address 203.0.113.250/30
```

6. Configure service options to be applied on the Multiservices interface.

```
[edit interfaces]
user@host# edit ms-11/1/0 services-options
```

NOTE: The values configured for the service options are shown for example only. You must configure and provision appropriate values as per the requirement.

7. Specify the open and close timeout periods in seconds for Transmission Control Protocol (TCP) session establishment.

```
[edit interfaces ms-11/1/0 services-options]
user@host# set open-timeout 4
user@host# set close-timeout 2
```

8. Specify the inactivity timeout periods in seconds for established TCP and non-TCP sessions.

```
[edit interfaces ms-11/1/0 services-options]
user@host# set inactivity-tcp-timeout 5
set inactivity-non-tcp-timeout 5
```

9. Specify the session lifetime in seconds globally for the Multiservices interface.

```
[edit interfaces ms-11/1/0 services-options]
user@host# set session-timeout 5
```

10. Specify the maximum number of keep-alive messages sent before a TCP session is allowed to time out.

```
[edit interfaces ms-11/1/0 services-options]  
user@host# set tcp-tickles 0
```

11. Configure a logical interface on the Multiservices interface.

```
[edit interfaces ms-11/1/0]  
user@host# edit unit 1
```

12. Configure the service domain to specify that the logical interface is used within the network.

```
[edit interfaces ms-11/1/0 unit 1]  
user@host# set service-domain inside
```

13. Configure the IPv4 address family on the logical interface.

```
[edit interfaces ms-11/1/0 unit 1]  
user@host# set family inet
```

14. Configure a second logical interface on the Multiservices interface.

```
[edit interfaces ms-11/1/0]  
user@host# edit unit 2
```

15. Configure the service domain to specify that the logical interface is used outside the network.

```
[edit interfaces ms-11/1/0 unit 2]  
user@host# set service-domain outside
```

16. Configure an output filter to redirect CPCD packets from the logical interface.

```
[edit interfaces ms-11/1/0 unit 2]  
user@host# set filter output FF_CPCD_REDIRECT_OUTPUT
```

17. Configure the IPv4 address family on the logical interface.

```
[edit interfaces ms-11/1/0 unit 2]
```

```
user@host# set family inet
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
root@host# show interfaces
ge-0/0/1 {
  unit 900 {
    description VLAN-REDIRECT;
    vlan-id 900;
  }
  family inet {
    filter {
      input FF_HTTP_REDIR_IN;
    }
    address 203.0.113.250/30;
  }
}
ms-11/1/0 {
  services-options {
    open-timeout 4;
    close-timeout 2;
    inactivity-tcp-timeout 5;
    inactivity-non-tcp-timeout 5;
    session-timeout 5;
    tcp-tickles 0;
  }
  unit 1 {
    family inet;
    service-domain inside;
  }
  unit 2 {
    family inet {
      filter {
        output FF_CPCD_REDIRECT_OUTPUT;
      }
    }
    service-domain outside;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Additional Routing Instance and Assigning Its Next-Hop Static Interfaces

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Configure a routing instance.

```
[edit routing-instances]
user@host# edit CPCD_REDIRECT
```

2. Configure a virtual router routing instance.

```
[edit routing-instances CPCD_REDIRECT]
user@host# set instance-type virtual-router
```

3. Configure the two previously defined multiservices interfaces for the routing instance.

```
[edit routing-instances CPCD_REDIRECT]
user@host# set interface ms-11/1/0.1
user@host# set interface ms-11/1/0.2
```

4. Configure static routing options.

```
[edit routing-instances CPCD_REDIRECT]
user@host# edit routing-options static
```

5. Assign the next-hop static interfaces to the routes and routing instance.

```
[edit routing-instances CPCD_REDIRECT routing-options static]
user@host# set route 0.0.0.0/0 next-hop ms-11/1/0.1
user@host# set route 203.0.113.0/24 next-hop ms-11/1/0.2
```

Results

From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
root@host# show routing-instances
CPCD_REDIRECT {
    instance-type virtual-router;
    interface ms-11/1/0.1;
    interface ms-11/1/0.2;
    routing-options {
        static {
            route 0.0.0.0/0 next-hop ms-11/1/0.1;
            route 203.0.113.0/24 next-hop ms-11/1/0.2;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Interface-Specific Filters to Direct HTTP Traffic

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Create a family for the service filter under the **[edit firewall]** hierarchy.

```
[edit firewall]
user@host# edit family inet
```

2. Create an interface-specific filter to redirect output traffic for CPCD.

```
[edit firewall family inet]
user@host# edit filter FF_CPCD_REDIRECT_OUTPUT
```

3. Specify that this is an interface-specific filter.

```
[edit firewall family inet filter FF_CPCD_REDIRECT_OUTPUT]
user@host# set interface-specific
```

4. Create a filter term for the interface-specific filter for the walled garden.

```
[edit firewall family inet filter FF_CPCD_REDIRECT_OUTPUT]
user@host# edit term One
```

- Specify both the action to count default traffic and the default routing instance.

```
[edit firewall family inet filter FF_CPCD_REDIRECT_OUTPUT interface-specific term One]
user@host# set then count back-to-default
set then routing-instance default
```

- Create a filter to redirect HTTP input traffic.

```
[edit firewall family inet]
user@host# edit filter FF_HTTP_REDIR_IN
```

- Specify that this is an interface-specific filter.

```
[edit firewall family inet filter FF_HTTP_REDIR_IN]
user@host# set interface-specific
```

- Create a filter term for the interface-specific filter for the walled garden.

```
[edit firewall family inet filter FF_HTTP_REDIR_IN]
user@host# edit term ACCEPTED_PREFIXES
```

- Specify the list of accepted prefixes as a match conditions for the walled garden's filter.

```
[edit firewall family inet filter FF_HTTP_REDIR_IN term ACCEPTED_PREFIXES]
user@host# set from prefix-list User-PRIVATE-Blocks-01
```

- Specify the action to take for all the matching HTTP traffic.

```
[edit firewall family inet filter FF_HTTP_REDIR_IN term ACCEPTED_PREFIXES]
user@host# set then next term
```

- Create a second filter term for the walled garden's filter.

```
[edit firewall family inet filter FF_HTTP_REDIR_IN interface-specific]
user@host# edit term HTTP
```

- Specify the protocol and destination port as match conditions for the walled garden's filter.

```
[edit firewall family inet filter FF_HTTP_REDIRECT_IN term HTTP]
user@host# set from protocol tcp
user@host# set from destination-port http
```

13. Specify the action to take for matching HTTP traffic destined to flow outside of the walled garden.

```
[edit firewall family inet filter filter FF_HTTP_REDIRECT_IN interface-specific term HTTP]
user@host# set then count HTTP
user@host# set then forwarding-class best-effort
user@host# set then routing-instance CPCD_REDIRECT
```

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]
root@host# show firewall
family inet {
  filter FF_CPCD_REDIRECT_OUTPUT {
    interface-specific;
    term One {
      then {
        count back-to-default;
        routing-instance default;
      }
    }
  }
  filter FF_HTTP_REDIRECT_IN {
    interface-specific;
    term ACCEPTED_PREFIXES {
      from {
        prefix-list {
          User-PRIVATE-Blocks-01;
        }
      }
      then next term;
    }
    term HTTP {
      from {
        protocol tcp;
        destination-port http;
      }
    }
  }
}
```

```

    }
    then {
        count http;
        forwarding-class best-effort;
        routing-instance CPCD_REDIRECT;
    }
}
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Policy Option and Statement to Use a Private Blocks Prefix List

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*.

1. Create a policy option and statement to use a private blocks prefix list under the **[edit policy-options]** hierarchy.

```

[edit policy-options]
user@host# set policy-statement User-PRIVATE-Blocks-01

```

2. Configure the source address for the private blocks prefix list.

```

[edit policy-options policy-statement User-PRIVATE-Blocks-01]
user@host# set 203.0.113.0/24

```

Results

From configuration mode, confirm your configuration by entering the **show policy-options** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
root@host# show policy-options
policy-statement User-PRIVATE-Blocks-01 {
    203.0.113.0/24;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying the Configured Service Set for CPCD Services | 549](#)
- [Verifying Details for a Configured HTTP Service Rule for a Walled Garden | 549](#)

To confirm that HTTP redirect services has been configured correctly within a service set, perform these tasks:

Verifying the Configured Service Set for CPCD Services

Purpose

Display the configured CPCD service set.

Action

From operational mode, enter the **show services captive-portal-content-delivery service-set http-redirect-sset detail** command.

```
user@host> show services captive-portal-content-delivery service-set http-redirect-sset detail
```

Service Set	Id	Profile	Compiled Rules
http-redirect-sset	1	http-redirect	1

Meaning

The output lists the service set configured for CPCD services.

Verifying Details for a Configured HTTP Service Rule for a Walled Garden

Purpose

Display details for a specific configured HTTP service rule for a walled garden.

Action

From operational mode, enter the **show services captive-portal-content-delivery rule redirect term REDIRECT** command.

```
user@host> show services captive-portal-content-delivery rule redirect term REDIRECT
```

```
Rule name: redirect
Rule match direction: input
Term name: term REDIRECT
Term action: redirect
Term action option: http://redirection-portal/redirection/
```

Meaning

The output lists rule and term details for a specific HTTP service rule configured for the walled garden.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview | 473](#)

Example: Configuring HTTP Redirect Services Using an Interface-Specific Filter and Attaching It to a Static Interface

7

PART

Configuring Subscriber Secure Policy

Configuring Subscriber Secure Policy Traffic Mirroring Overview | **553**

Configuring RADIUS-Initiated Subscriber Secure Policy Traffic Mirroring | **557**

Configuring DTCP-Initiated Subscriber Secure Policy Traffic Mirroring | **579**

Configuring DTCP Messages Used for DTCP-Initiated Subscriber Secure Policy Mirroring | **603**

Configuring Subscriber Secure Policy Support for IPv4 Multicast Traffic | **621**

Configuring Intercept-Related Information for Subscriber Secure Policy | **623**

Configuring Subscriber Secure Policy Traffic Mirroring Overview

IN THIS CHAPTER

- [Subscriber Secure Policy Overview | 553](#)

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 (SSP) 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, the subscriber's incoming and outgoing traffic are both 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.

NOTE: This feature requires a license. To understand more about Subscriber Access Licensing, see, [Subscriber Access Licensing Overview](#). Please refer to the [Juniper Licensing Guide](#) for general information about License Management. Please refer to the product Data Sheets at [MX Series 5G Universal Routing Platform](#) for details, or contact your Juniper Account Team or Juniper Partner.

Support for Intercepting Both Layer 2 and Layer 3 Datagrams

When DTCP- or RADIUS-initiated SSP intercepts traffic on logical subscriber interfaces and VLAN subscriber interfaces, it sends both Layer 2 and Layer 3 datagrams to the mediation device. 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.

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*.

Starting in Junos OS Release 16.1R1, you must configure the target parameters for mediation devices so that the SNMPv3 traps are sent with privacy (encrypted). Targets without privacy configured cannot receive the notifications.

In earlier releases, you can configure target parameters without privacy, allowing unencrypted notifications to be sent to the mediation devices. You also cannot restrict the traps to specific targets.

Support for L2TP Subscribers

Both DTCP-initiated and RADIUS-initiated SSP can be applied to Point-to-Point Protocol (PPP) subscribers whose traffic is tunneled with Layer 2 Tunneling Protocol (L2TP). DTCP SSP supports subscribers only at the L2TP network server (LNS), whereas RADIUS-initiated SSP supports subscribers at the L2TP access concentrator (LAC) or the LNS.

At the LAC, 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. 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.

At the LNS, both subscriber ingress traffic (from the 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.

There is no specific L2TP SSP configuration.

Junos OS Service for Subscriber Secure Policy Traffic Mirroring

Subscriber secure policy mirroring requires the use of the radius-flow-tap service, configured at the **[edit services radius-flow-tap]** hierarchy level. This service is used only for subscriber secure policy mirroring and only on MX Series routers.

There are other Junos OS services with similar names, but they are not used for subscriber secure policy mirroring:

- The flow-tap service, configured at the **[edit services flow-tap]** hierarchy level, is an older Junos OS service for packet mirroring. This service uses Dynamic Tasking Control Protocol (DTCP) requests from mediation devices to intercept IPv4 packets in an active flow monitoring station (router). The router uses DTCP to send a copy of packets that match filter criteria to one or more content destinations. The flow-tap service is supported only on M Series and T Series routers using Adaptive Services PICs. For information about the flow-tap service, see *Understanding Flow-Tap Architecture*.
- The FlowTapLite service is a lightweight version of the flow-tap service for packet mirroring. It is also configured at the **[edit services flow-tap]** hierarchy level. The FlowTapLite service resides on the Packet Forwarding Engine rather than a line card. The intercepted packets are sent to a tunnel logical interface (vt-) for encapsulation, so you must allocate and assign tunnel interfaces for the service. It is supported on MX Series routers and on M320 routers with Enhanced III Flexible PIC Concentrators (FPCs). You cannot run FlowTapLite and the flow-tap service on the same router concurrently. For information about FlowTapLite, see *Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs*.

Protection of SSP Data when a Core Error is Generated

When the authd, bbe-smgd, or dfcd processes generate a core error, the core dump file contains information related to whatever the process is involved with, including SSP. The error files contain SSP information that might identify the subscriber whose traffic is mirrored or the mediation device that receives the mirrored traffic. For example, the files include information such as the source and destination IP address for the mediation device, device ports, and intercept ID.

Starting in Junos OS Release 18.4R1, SSP-related information is automatically encrypted in core dump files to prevent this information from being seen by unauthorized persons in the event of a core error. Encryption is enabled by default; no configuration is required or possible. The dfcd core error files may contain traffic mirroring information that does not identify subscribers or devices; this information is not masked. FlowTapLite information is not masked.

NOTE: Information related to SSP is not encrypted when it is in a transient state; for example, if the core error occurs when the data has been received from a RADIUS or DTCP server, but is not yet encrypted.

Subscriber Secure Policy Licensing Requirements

To enable and use subscriber secure policy, you must install and properly configure the Subscriber Secure Policy license.

Release History Table

Release	Description
18.4R1	Starting in Junos OS Release 18.4R1, SSP-related information is automatically encrypted in core dump files to prevent this information from being seen by unauthorized persons in the event of a core error.

RELATED DOCUMENTATION

RADIUS-Initiated Subscriber Secure Policy Overview	557
DTCP-Initiated Subscriber Secure Policy Overview	579
Intercept-Related Events Transmitted to the Mediation Device	623

Configuring RADIUS-Initiated Subscriber Secure Policy Traffic Mirroring

IN THIS CHAPTER

- [RADIUS-Initiated Subscriber Secure Policy Overview | 557](#)
- [Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS | 558](#)
- [RADIUS-Initiated Traffic Mirroring Interfaces | 560](#)
- [RADIUS-Initiated Traffic Mirroring Process at Subscriber Login | 563](#)
- [RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers | 564](#)
- [RADIUS Attributes Used for Subscriber Secure Policy | 566](#)
- [Using the Packet Header to Track Subscribers on the Mediation Device | 567](#)
- [Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)
- [Guidelines for Configuring Subscriber Secure Policy Mirroring | 573](#)
- [Configuring Support for Subscriber Secure Policy Mirroring | 574](#)
- [Configuring RADIUS Server Support for Subscriber Secure Policy Mirroring | 577](#)
- [Terminating RADIUS-Initiated Subscriber Traffic Mirroring | 578](#)

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 | 558](#)
- [Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS

Figure 11 on page 558 shows the architecture of the RADIUS-initiated subscriber secure policy mirroring environment.

Figure 11: RADIUS-Initiated Subscriber Secure Policy Architecture

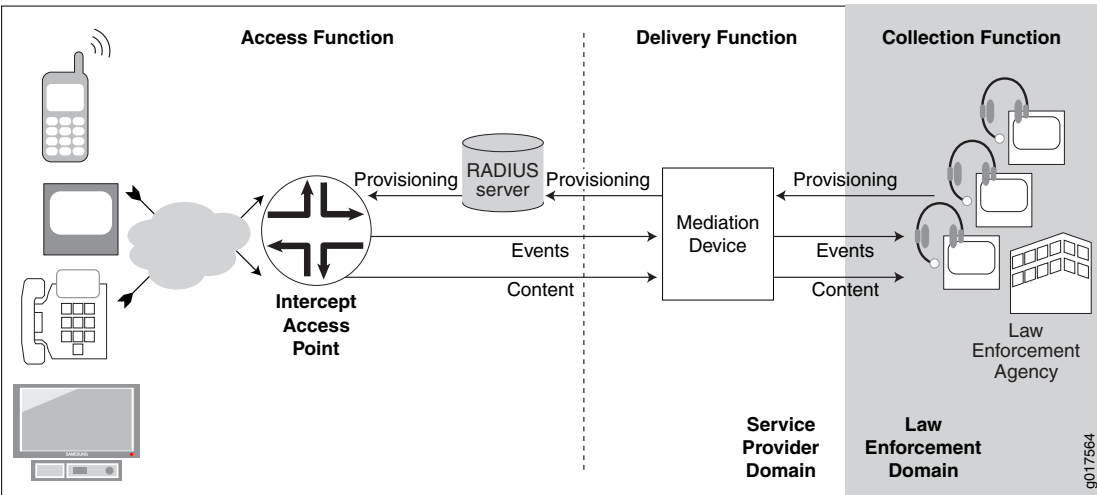


Table 39 on page 559 describes the functions and components of a RADIUS-initiated subscriber secure policy traffic mirroring environment.

Table 39: RADIUS-Initiated Subscriber Secure Policy Functions and Components

Function or Component	Description
Collection function	<p>The collection function is responsible for collecting intercepted content and identifying information from the delivery function.</p> <p>The collection function is the responsibility of the law enforcement agency (LEA).</p>
Delivery function	<p>The delivery function delivers information that it receives from the access function to the collection function.</p> <p>The delivery function is performed by the mediation device.</p>
Access function	<p>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.</p> <p>The access function is the responsibility of intercept access points (IAPs).</p>
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	<p>The mediation device receives provisioning information from the LEA, and it uses the information to send provisioning information to the RADIUS server.</p> <p>The mediation device also receives intercept-related events and intercepted content from the router, and delivers the events and intercepted content to the LEA.</p>
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.

Table 39: RADIUS-Initiated Subscriber Secure Policy Functions and Components *(continued)*

Function or Component	Description
IAP	<p>Intercept access point. In a subscriber access network the Juniper Networks router is the IAP.</p> <p>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.</p> <p>The IAP also sends intercept-related events to the mediation device using SNMP traps.</p>

RELATED DOCUMENTATION

RADIUS-Initiated Subscriber Secure Policy Overview 557
RADIUS-Initiated Traffic Mirroring Interfaces 560
RADIUS-Initiated Traffic Mirroring Process at Subscriber Login 563
RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers 564

RADIUS-Initiated Traffic Mirroring Interfaces

Figure 12 on page 561 shows the interfaces involved in RADIUS-initiated secure subscriber policy traffic mirroring.

Figure 12: RADIUS-Initiated Traffic Mirroring Interfaces

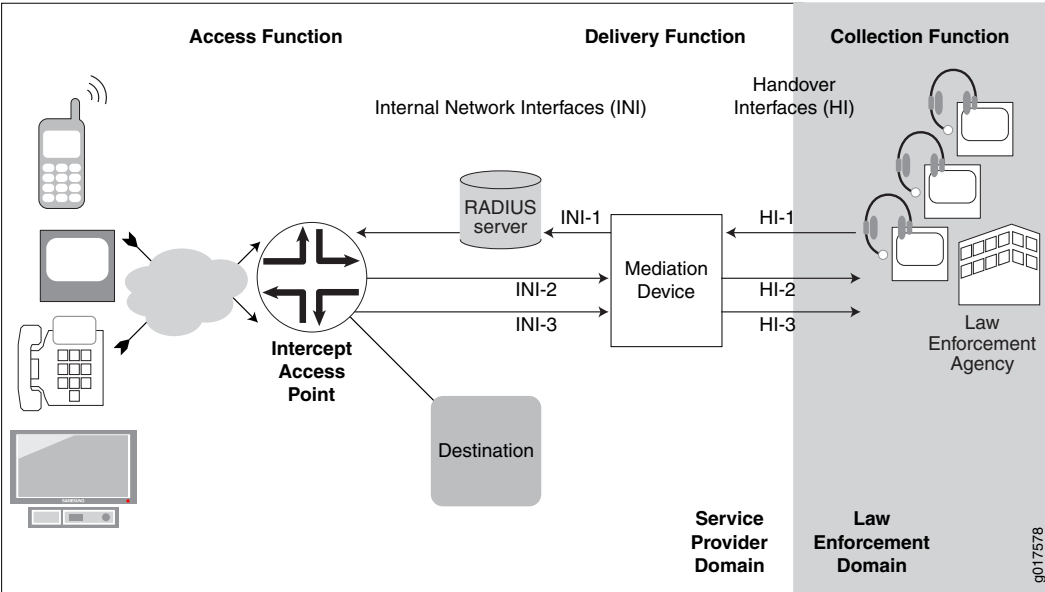


Table 40 on page 561 describes the interfaces involved in RADIUS-initiated secure subscriber policy traffic mirroring.

Table 40: 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.

Table 40: RADIUS-Initiated Traffic Mirroring Interfaces (*continued*)

Interface	Description
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.

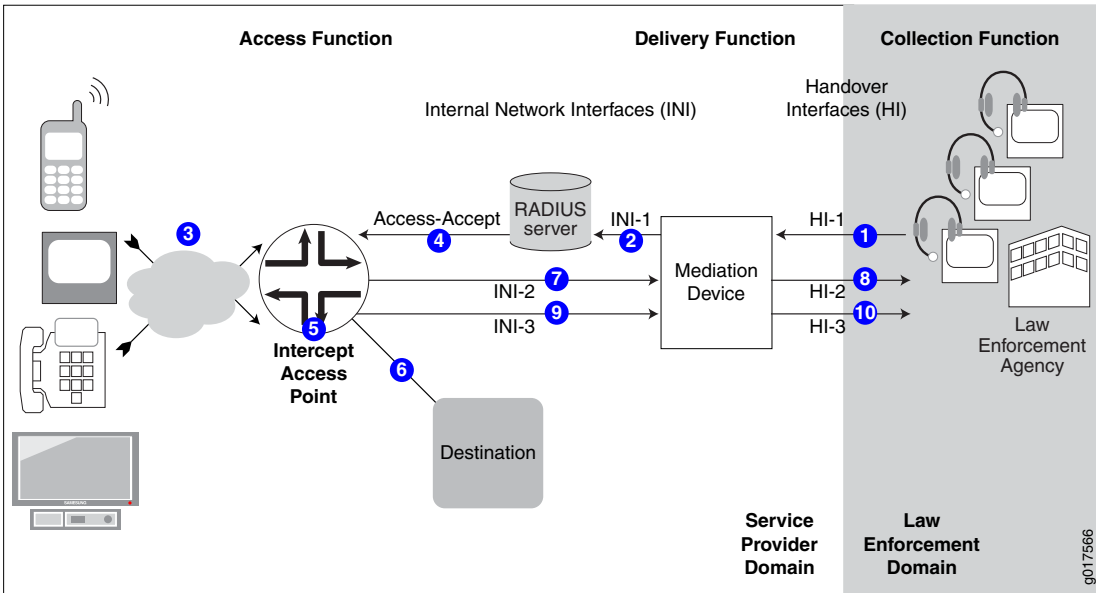
RELATED DOCUMENTATION

Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS 558
RADIUS-Initiated Traffic Mirroring Process at Subscriber Login 563
RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers 564

RADIUS-Initiated Traffic Mirroring Process at Subscriber Login

Figure 13 on page 563 shows the process for a RADIUS-initiated subscriber mirroring operation that is initiated when the mirrored subscriber logs in.

Figure 13: 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 message 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 | 558](#)

[RADIUS-Initiated Traffic Mirroring Interfaces | 560](#)

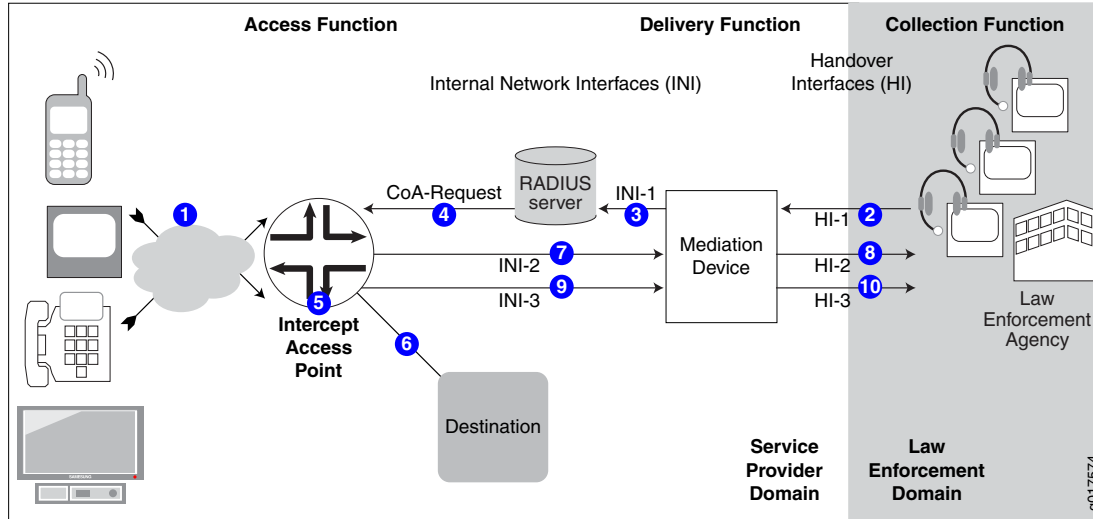
[RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers | 564](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

RADIUS-Initiated Traffic Mirroring Process for Logged-In Subscribers

[Figure 14 on page 565](#) shows the process for a RADIUS-initiated subscriber mirroring operation that is initiated after the subscriber has logged in.

Figure 14: 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).

2–The LEA sends provisioning information for a subscriber whose traffic is to be mirrored over the HI-1 interface to the mediation device.

3–The mediation device sends the provisioning information over the INI-1 interface to the RADIUS server.

4–The RADIUS server sends a CoA message containing the mirroring-related RADIUS attributes and VSAs to the IAP (the router).

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.

6–The IAP sends the original subscriber traffic to its intended destination.

7–As subscriber-related events occur, the IAP sends the events in SNMP traps over the INI-2 interface to the mediation device.

8–The mediation device provides events over the HI-2 interface to the LEA.

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 message from the RADIUS server.

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 | 558](#)

[RADIUS-Initiated Traffic Mirroring Interfaces | 560](#)

[RADIUS-Initiated Traffic Mirroring Process at Subscriber Login | 563](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

RADIUS Attributes Used for Subscriber Secure Policy

Table 41 on page 566 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 41: RADIUS-Based Mirroring Attributes

Attribute Number	Attribute Name	Description	Value
[26-58]	LI-Action	Traffic mirroring action	Salt-encrypted integer <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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; otherwise, the action fails.

If a subscriber is already logged in, [Table 42 on page 567](#) 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).

Table 42: 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, traffic mirroring fails when the DHCP session is authenticated and activated.

RELATED DOCUMENTATION

[RADIUS-Initiated Subscriber Secure Policy Overview | 557](#)

[Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS | 558](#)

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 15 on page 568](#) is the mirrored packet header and payload that the router sends to the mediation device.

Figure 15: Mirrored Packet Header and Payload




Table 43 on page 568 describes the fields in the packet header of mirrored packets.

Table 43: Mirrored Packet Header and Payload Field Descriptions For the Mediation Device

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

Table 43: Mirrored Packet Header and Payload Field Descriptions For the Mediation Device (continued)

Field	Value	Length (Bits)
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 (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 570 for details	30

Table 43: Mirrored Packet Header and Payload Field Descriptions For the Mediation Device (*continued*)

Field	Value	Length (Bits)
Session-ID	See “Format of the Mirror Header Values Used to Track Subscribers and Subscriber Sessions” on page 570 for details	32

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 16 on page 571](#):

- **V = 1**

- Intercept ID = 0x10

Figure 16: 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.

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 00000300000000090 in VSA 26-59 configures the following fields in the mirror header, as shown in [Figure 17 on page 571](#):

- V = 0
- Intercept-ID = 0x300
- Session-ID = 0x90

Figure 17: 8-Byte Format of VSA 26-59



RELATED DOCUMENTATION

[RADIUS-Initiated Subscriber Secure Policy Overview | 557](#)

[Subscriber Secure Policy Traffic Mirroring Architecture Using RADIUS | 558](#)

Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview

Before you configure subscriber secure policy traffic mirroring, note the following:

- Subscriber secure policy mirroring runs on the radius-flow-tap service infrastructure. To configure the subscriber secure policy service, you must have the same privileges that are required to configure the radius-flow-tap service.
- The subscriber secure policy feature requires some system resources while mirroring, encrypting, and sending traffic to the mediation device. For example, you might elect to use a 10-Gigabit Ethernet interface for the tunnel to the mediation device if you expect the amount of traffic you plan to mirror to approach 1 Gbps of actual user data.

To configure the subscriber secure policy service:

1. Configure radius-flow-tap service support for secure subscriber policy. This support includes optional forwarding-class information that the subscriber secure policy service uses to send mirrored traffic to the content destination device.

See [“Configuring Support for Subscriber Secure Policy Mirroring” on page 574](#).

2. Configure an access profile that specifies the RADIUS-related support for subscriber secure policy on the router, including a list of one or more RADIUS authentication servers. The router uses the list of specified servers for both authentication and dynamic request operations. You must also configure the RADIUS dynamic request feature, which provides the CoA message support used in-session traffic mirroring.

See [“Configuring RADIUS Server Support for Subscriber Secure Policy Mirroring” on page 577](#).

3. Ensure that the following support is also configured:

- The RADIUS record of the mirrored subscriber must include the RADIUS attributes and VSAs required for subscriber secure policy mirroring. See [“RADIUS Attributes Used for Subscriber Secure Policy” on page 566](#) 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.

4. (Optional) Enable the mirroring of IPv4 multicast traffic on the router.

See [“Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic” on page 622](#).

5. (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 625](#).

You can terminate an active subscriber mirroring session at any time.

See [“Terminating RADIUS-Initiated Subscriber Traffic Mirroring”](#) on page 578.

RELATED DOCUMENTATION

[RADIUS Attributes Used for Subscriber Secure Policy | 566](#)

[Guidelines for Configuring Subscriber Secure Policy Mirroring | 573](#)

[Intercept-Related Events Transmitted to the Mediation Device | 623](#)

[Terminating RADIUS-Initiated Subscriber Traffic Mirroring | 578](#)

Guidelines for Configuring Subscriber Secure Policy Mirroring

The subscriber secure policy service uses the radius-flow-tap service infrastructure. Consider the following guidelines when you configure subscriber secure policy mirroring:

When configuring subscriber secure policy mirroring, consider the following guidelines regarding the relationship between the radius-flow-tap service and the FlowTapLite service on MX Series tunnel interfaces (FlowTapLite):

- Starting in Junos OS Release 17.3R1, the radius-flow-tap service can run concurrently on the same router with the FlowTapLite service. The FlowTapLite service is a version of the flow-tap service (**[edit services flow-tap]**) that is configured only on tunnel interfaces on MX Series routers and is not used for subscriber secure policy mirroring.

In earlier releases, the radius-flow-tap and FlowTapLite services cannot run concurrently on an MX Series router, preventing you from running FlowTapLite monitoring and subscriber secure policy mirroring at the same time.

- You can configure one instance of the radius-flow-tap service on the router. Subscriber secure policy RADIUS-initiated mirroring and Dynamic Tasking Control Protocol (DTCP)-initiated mirroring both use the radius-flow-tap service.
- If you delete the radius-flow-tap service, new subscribers are not monitored. Existing subscribers that already have subscriber secure policy attached are not affected when you delete the service configuration.
- You can retain DTCP-initiated mirroring but prevent RADIUS-initiated mirroring from being enabled by including the **[edit system services dtcp-only]** statement, if you do so before any RADIUS-initiated mirroring is attached to a subscriber. Subsequently, RADIUS requests to initiate mirroring are rejected; only DTCP-initiated mirroring and FlowTapLite are allowed. Existing RADIUS-initiated mirroring services are not affected.
- Starting in Junos OS Release 16.1R1, you must configure the target parameters for mediation devices so that the SNMPv3 traps are sent with privacy (encrypted). Targets without privacy configured cannot receive the notifications. In earlier releases, you can configure target parameters without privacy, allowing

unencrypted notifications to be sent to the mediation devices. You must also explicitly configure a list of trap targets with the `[edit services radius-flow-tap snmp notify-targets]` statement.

Release History Table

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, the radius-flow-tap service can run concurrently on the same router with the FlowTapLite service.
16.1R1	Starting in Junos OS Release 16.1R1, you must configure the target parameters for mediation devices so that the SNMPv3 traps are sent with privacy (encrypted).

RELATED DOCUMENTATION

[Subscriber Secure Policy Overview | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

[Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview | 590](#)

[Configuring Support for Subscriber Secure Policy Mirroring | 574](#)

[Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring | 597](#)

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. Specify how the mirrored packets are forwarded to the mediation device.

- To mirror interfaces created by extensible subscriber services manager (ESSM), assign the virtual tunnel interfaces for the radius-flow-tap service.

```
[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.

- To mirror flow-based interfaces, specify the logical system and routing instance for the radius-flow-tap service.

```
[edit services radius-flow-tap]
user@host# set logical-system LS1 routing-instance RI1
```

You can specify a logical system and routing instance, or a routing instance without a logical system. If you do not specify a logical system, the router uses logical system **default**. If you do not specify either a logical system or routing instance, the router uses logical system **default** and routing instance **default**.

BEST PRACTICE: Configure a routing instance to prevent a spoofed mediation device address from diverting traffic away from the device. When the mirrored customer flows are in the same routing instance as the mediation device, a malicious user might hijack the mediation device's route advertisement. By advertising a next hop to the hijacker's network instead of to the device, the mirrored flows are captured and never reach the mediation device.

If you configure the mirrored traffic to be forwarded to the mediation device by means of a routing instance, then the traffic is separated from the Internet. An external user is then unable to divert the mirrored traffic to the user's network.

NOTE: The **interfaces** statement applies only to ESSM-created interfaces and is ignored for flow-based interfaces. Similarly, the LS:RI configuration applies only to flow-based interfaces.

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
```

5. (Optional) Specify the subscriber secure policy that determines what traffic, if any, is not sent to the mediation device.

```
[edit services radius-flow-tap]
user@host# set policy policy-name
```

NOTE: You can add or change a subscriber secure policy any time, but a changed policy does not apply to a currently enabled policy. To change a policy:

- Send a DTCP DELETE message to remove the current policy.
- Modify the configuration with the new version of the policy.
- Send a DTCP ADD message to add the policy.
- Send a DTCP ENABLE message to enable the policy.

6. (Optional) Specify the IP address for one or more target mediation devices to receive SNMPv3 trap notifications. Each target address must be configured separately.

```
[edit services radius-flow-tap]
user@host# set snmp notify-targets ip-address
```

NOTE: You must also configure SNMP so that only encrypted notifications are sent to target devices. Targets without privacy configured cannot receive the notifications. For information about the SNMP configuration for subscriber secure policy, see [“Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring”](#) on page 625.

RELATED DOCUMENTATION

[Subscriber Secure Policy Overview | 553](#)[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)[Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview | 590](#)[Guidelines for Configuring Subscriber Secure Policy Mirroring | 573](#)

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

[RADIUS Attributes Used for Subscriber Secure Policy | 566](#)

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 | 557](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

Configuring DTCP-Initiated Subscriber Secure Policy Traffic Mirroring

IN THIS CHAPTER

- [DTCP-Initiated Subscriber Secure Policy Overview | 579](#)
- [Subscriber Secure Policy Traffic Mirroring Architecture Using DTCP | 580](#)
- [DTCP-Initiated Traffic Mirroring Interfaces | 582](#)
- [DTCP-Initiated Traffic Mirroring Process | 584](#)
- [DTCP Messages Used for Subscriber Secure Policy | 585](#)
- [Packet Header for Mirrored Traffic Sent to Mediation Device | 586](#)
- [Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview | 590](#)
- [Guidelines for Configuring Subscriber Secure Policy Mirroring | 591](#)
- [Configuring Support for Subscriber Secure Policy Mirroring | 592](#)
- [Configuring the Mediation Device as a User on the Router | 595](#)
- [Configuring a DTCP-over-SSH Connection to the Mediation Device | 596](#)
- [Configuring the Mediation Device to Provision Traffic Mirroring | 597](#)
- [Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring | 597](#)
- [Example: Configuring Traffic That Is Mirrored Using DTCP-Initiated Subscriber Secure Policy | 598](#)
- [Terminating DTCP-Initiated Subscriber Traffic Mirroring Sessions | 601](#)

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 sent from the mediation device 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 | 580](#)
- [Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview | 590](#)

Subscriber Secure Policy Traffic Mirroring Architecture Using DTCP

Figure 18 on page 580 shows the architecture of the DTCP-initiated subscriber secure policy mirroring environment.

Figure 18: DTCP-Initiated Subscriber Secure Policy Architecture




Table 44 on page 580 describes the functions and components of a DTCP-initiated subscriber secure policy traffic mirroring environment.

Table 44: DTCP-Initiated Subscriber Secure Policy Functions and Components

Function or Component	Description
Collection function	<p>The collection function is responsible for collecting intercepted content and identifying information from the delivery function.</p> <p>The collection function is the responsibility of the law-enforcement agency (LEA).</p>

Table 44: DTCP-Initiated Subscriber Secure Policy Functions and Components (*continued*)

Function or Component	Description
Delivery function	<p>The delivery function delivers information that it receives from the access function to the collection function.</p> <p>The delivery function is performed by the mediation device.</p>
Access function	<p>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.</p> <p>The access function is performed by intercept access points (IAPs).</p>
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	<p>The mediation device receives provisioning information from the LEA, and it uses the information to send provisioning information to the IAP (the router).</p> <p>The mediation device also receives intercept-related events and intercepted content from the router, and delivers the events and content to the LEA.</p>
IAP	<p>Intercept access point. In a subscriber access network the Juniper Networks router is the IAP.</p> <p>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.</p> <p>The IAP also sends intercept-related events to the mediation device using SNMP traps.</p>

RELATED DOCUMENTATION

[DTCP-Initiated Subscriber Secure Policy Overview | 579](#)
[DTCP-Initiated Traffic Mirroring Interfaces | 582](#)
[DTCP-Initiated Traffic Mirroring Process | 584](#)

DTCP-Initiated Traffic Mirroring Interfaces

Figure 19 on page 582 shows the interfaces involved in DTCP-initiated secure subscriber policy traffic mirroring.

Figure 19: DTCP-Initiated Traffic Mirroring Interfaces




Table 45 on page 582 describes the interfaces involved in DTCP-initiated secure subscriber policy traffic mirroring.

Table 45: 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.

Table 45: DTCP-Initiated Traffic Mirroring Interfaces *(continued)*

Interface	Description
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 580
DTCP-Initiated Traffic Mirroring Process 584

DTCP-Initiated Traffic Mirroring Process

Figure 20 on page 584 shows the process for a DTCP-initiated subscriber mirroring operation.

Figure 20: 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 | 580

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 four types of DTCP messages supported for radius-flow-tap services:

- **ADD**—Triggers mirroring of subscriber secure policy sessions. You include attributes that trigger the router to begin mirroring a subscriber session. In addition to one or more attributes that trigger the router to begin traffic mirroring, 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.
- **DELETE**—Removes a subscriber mirroring trigger or can be used to remove all mirroring.
- **ENABLE**—Triggers a drop policy on the router if one does not already exist from a prior DTCP ADD or DTCP ENABLE message.
- **LIST**—Requests information about sessions that are currently being mirrored. This information is returned in a LIST response.

NOTE: Consult the documentation for your mediation device to learn how to configure DTCP messages on the device.

RELATED DOCUMENTATION

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.

The packet header includes the Session ID that Junos assigns to the subscriber session. The mediation device can use the ID to identify the session of the mirrored subscriber. The mediation device can use the Session ID along with the Intercept ID to track a subscriber across multiple login and logout events. The Intercept ID is constant, but the Session ID changes with each new session for a subscriber.

Figure 21 on page 586 is the mirrored packet header that the router sends to the mediation device.

Figure 21: Mirrored Packet Header and Payload

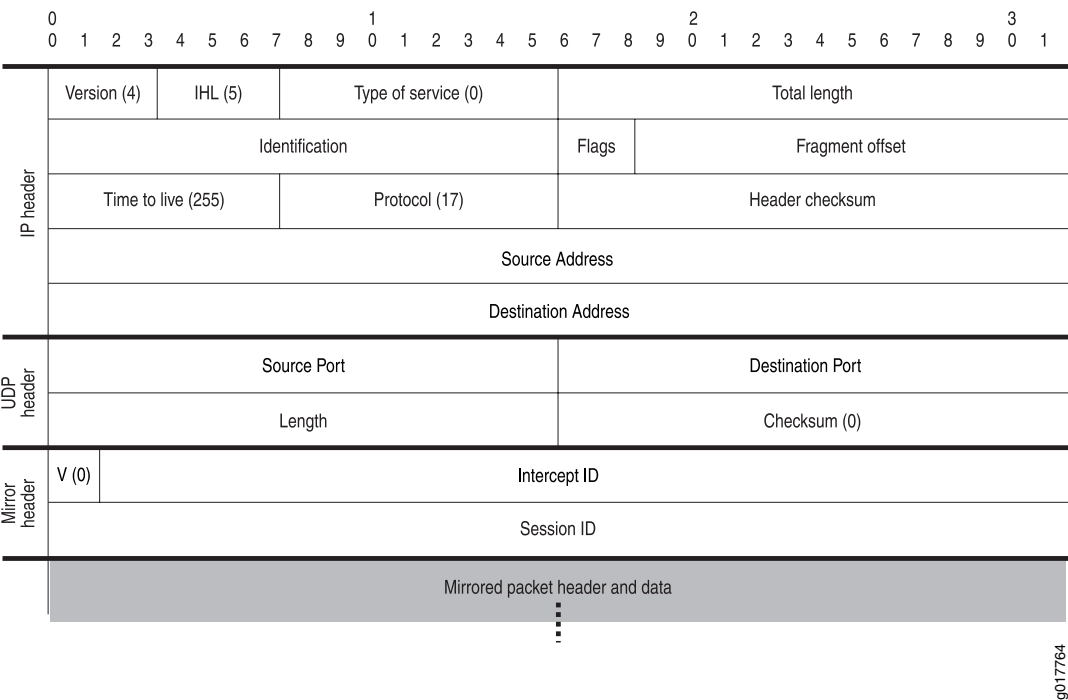


Table 46 on page 586 describes the fields in the packet header of mirrored packets.

Table 46: Packet Header Field Descriptions

Field	Value	Length (Bits)
IP Header		
Version	4	4
IHL	5	4

Table 46: Packet Header Field Descriptions (*continued*)

Field	Value	Length (Bits)
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
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

Table 46: Packet Header Field Descriptions (*continued*)

Field	Value	Length (Bits)
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
Session ID	Subscriber session ID assigned by the router. See “Format of the Mirror Header Values Used to Track Subscribers and Subscriber Sessions” on page 588	32

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. There are three mirror header attributes in the packet header:

- V (mirror header value)— For DTCP, this value is always set to 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.

Manually Setting the Session-ID and Intercept ID in Packet Headers


You can use the DTCP ADD command to manually specify the Session ID value (**X-Act-Sess-Id**) and the Intercept ID value (**X-MD-Intercept-Id**) placed in the headers sent to the mediation device. You configure the values in an 8-byte format. To do so:

- Configure the first two most significant bits to a value of 0, which indicates two words.
- Configure the remaining 30 bits of the first word to form the Intercept ID field.
- Configure the second word to form the Session-ID field.

You cannot change the order of these two words.

Figure 22 on page 589 shows an example of the mirror header:

Figure 22: Mirror Header Format



For example, a value of 0000030000000090 configures the following fields in the mirror header: :

- V = 0
- Intercept-ID = 0x300
- Session-ID = 0x90

RELATED DOCUMENTATION

[DTCP-Initiated Subscriber Secure Policy Overview | 579](#)

[ADD \(DTCP\) | 604](#)

[Example: Using DTCP Messages to Trigger, Verify, and Remove Traffic Mirroring for Subscribers | 615](#)

Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview

Before you configure subscriber secure policy traffic mirroring, note the following:

- Subscriber secure policy mirroring runs on the radius-flow-tap service infrastructure. To configure the subscriber secure policy service, you need the same privileges that are required to configure the radius-flow-tap service.
- The subscriber secure policy feature requires some system resources while mirroring, encrypting, and sending traffic to the mediation device. For example, you might elect to use a 10-Gigabit Ethernet interface for the tunnel and mediation device if you expect the amount of traffic you plan to mirror to approach 1 Gbps of actual user data.

To configure DTCP-initiated subscriber secure policy service:

1. Configure the radius-flow-tap service support for secure subscriber policy. This support includes configuring the tunnels and optional forwarding-class information that the subscriber secure policy service uses to send mirrored traffic to the content destination device.
See [“Configuring Support for Subscriber Secure Policy Mirroring” on page 574](#).
2. Configure the mediation device as a user on the router. This user account allows the router to receive DTCP messages from the mediation device.
See [“Configuring the Mediation Device as a User on the Router” on page 595](#).
3. Configure the mediation device to provision traffic mirroring on the router.
See [“Configuring the Mediation Device to Provision Traffic Mirroring” on page 597](#).
4. Configure a DTCP-over-SSH connection to the mediation device.
See [“Configuring a DTCP-over-SSH Connection to the Mediation Device” on page 596](#).
5. (Optional) Enable mirroring of IPv4 multicast traffic on the router.
See [“Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic” on page 622](#).
6. Configure SNMPv3 trap support to report mirroring information to an external device.
See [“Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring” on page 625](#).

You can terminate an active subscriber mirroring session at any time.

See [“Terminating DTCP-Initiated Subscriber Traffic Mirroring Sessions” on page 601](#).

RELATED DOCUMENTATION

[DTCP-Initiated Subscriber Secure Policy Overview | 579](#)

[Intercept-Related Events Transmitted to the Mediation Device | 623](#)

Guidelines for Configuring Subscriber Secure Policy Mirroring

The subscriber secure policy service uses the radius-flow-tap service infrastructure. Consider the following guidelines when you configure subscriber secure policy mirroring:

When configuring subscriber secure policy mirroring, consider the following guidelines regarding the relationship between the radius-flow-tap service and the FlowTapLite service on MX Series tunnel interfaces (FlowTapLite):

- Starting in Junos OS Release 17.3R1, the radius-flow-tap service can run concurrently on the same router with the FlowTapLite service. The FlowTapLite service is a version of the flow-tap service (**[edit services flow-tap]**) that is configured only on tunnel interfaces on MX Series routers and is not used for subscriber secure policy mirroring.

In earlier releases, the radius-flow-tap and FlowTapLite services cannot run concurrently on an MX Series router, preventing you from running FlowTapLite monitoring and subscriber secure policy mirroring at the same time.

- You can configure one instance of the radius-flow-tap service on the router. Subscriber secure policy RADIUS-initiated mirroring and Dynamic Tasking Control Protocol (DTCP)-initiated mirroring both use the radius-flow-tap service.
- If you delete the radius-flow-tap service, new subscribers are not monitored. Existing subscribers that already have subscriber secure policy attached are not affected when you delete the service configuration.
- You can retain DTCP-initiated mirroring but prevent RADIUS-initiated mirroring from being enabled by including the **[edit system services dtcp-only]** statement, if you do so before any RADIUS-initiated mirroring is attached to a subscriber. Subsequently, RADIUS requests to initiate mirroring are rejected; only DTCP-initiated mirroring and FlowTapLite are allowed. Existing RADIUS-initiated mirroring services are not affected.
- Starting in Junos OS Release 16.1R1, you must configure the target parameters for mediation devices so that the SNMPv3 traps are sent with privacy (encrypted). Targets without privacy configured cannot receive the notifications. In earlier releases, you can configure target parameters without privacy, allowing unencrypted notifications to be sent to the mediation devices. You must also explicitly configure a list of trap targets with the **[edit services radius-flow-tap snmp notify-targets]** statement.

Release History Table

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, the radius-flow-tap service can run concurrently on the same router with the FlowTapLite service.
16.1R1	Starting in Junos OS Release 16.1R1, you must configure the target parameters for mediation devices so that the SNMPv3 traps are sent with privacy (encrypted).

RELATED DOCUMENTATION

[Subscriber Secure Policy Overview | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

[Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview | 590](#)

[Configuring Support for Subscriber Secure Policy Mirroring | 574](#)

[Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring | 597](#)

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. Specify how the mirrored packets are forwarded to the mediation device.

- To mirror interfaces created by extensible subscriber services manager (ESSM), assign the virtual tunnel interfaces for the radius-flow-tap service.

```
[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.

- To mirror flow-based interfaces, specify the logical system and routing instance for the radius-flow-tap service.

```
[edit services radius-flow-tap]
user@host# set logical-system LS1 routing-instance RI1
```

You can specify a logical system and routing instance, or a routing instance without a logical system. If you do not specify a logical system, the router uses logical system **default**. If you do not specify either a logical system or routing instance, the router uses logical system **default** and routing instance **default**.

BEST PRACTICE: Configure a routing instance to prevent a spoofed mediation device address from diverting traffic away from the device. When the mirrored customer flows are in the same routing instance as the mediation device, a malicious user might hijack the mediation device's route advertisement. By advertising a next hop to the hijacker's network instead of to the device, the mirrored flows are captured and never reach the mediation device.

If you configure the mirrored traffic to be forwarded to the mediation device by means of a routing instance, then the traffic is separated from the Internet. An external user is then unable to divert the mirrored traffic to the user's network.

NOTE: The **interfaces** statement applies only to ESSM-created interfaces and is ignored for flow-based interfaces. Similarly, the LS:RI configuration applies only to flow-based interfaces.

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
```

5. (Optional) Specify the subscriber secure policy that determines what traffic, if any, is not sent to the mediation device.

```
[edit services radius-flow-tap]
user@host# set policy policy-name
```

NOTE: You can add or change a subscriber secure policy any time, but a changed policy does not apply to a currently enabled policy. To change a policy:

- Send a DTCP DELETE message to remove the current policy.
- Modify the configuration with the new version of the policy.
- Send a DTCP ADD message to add the policy.
- Send a DTCP ENABLE message to enable the policy.

6. (Optional) Specify the IP address for one or more target mediation devices to receive SNMPv3 trap notifications. Each target address must be configured separately.

```
[edit services radius-flow-tap]
user@host# set snmp notify-targets ip-address
```

NOTE: You must also configure SNMP so that only encrypted notifications are sent to target devices. Targets without privacy configured cannot receive the notifications. For information about the SNMP configuration for subscriber secure policy, see [“Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring”](#) on page 625.

RELATED DOCUMENTATION

[Subscriber Secure Policy Overview](#) | 553

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

[Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview | 590](#)

[Guidelines for Configuring Subscriber Secure Policy Mirroring | 573](#)

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 uiduid-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 encrypted-password
```

Configuring a DTCP-over-SSH Connection to the Mediation Device

DTCP-initiated subscriber secure policy requires a DTCP-over-SSH connection for the radius-flow-tap service. This connection is used to send provisioning information from the mediation device to the router.

NOTE: DTCP-over-SSH connections are used for flow-tap, FlowTapLite, and radius-flow-tap services.

To configure the DTCP-over-SSH connection to support subscriber secure policy mirroring:

1. Access the **flow-tap-dtcp** hierarchy level.

```
[edit system services]
user@host# edit flow-tap-dtcp
```

NOTE:

2. Enable SSH support for DTCP.

```
[edit system services flow-tap-dtcp]
user@host# set ssh
```

3. (Optional) Configure the 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](#) | 590

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** message.
- To remove an existing traffic-mirroring trigger, use the **DELETE** message.
- To configure attributes to trigger a drop policy on the router (if one does not already exist), use the **ENABLE** message.
- To show existing traffic-mirroring triggers, use the **LIST** message.

NOTE: Consult the documentation for your mediation device to learn how to configure DTCP messages on the device.

For an example of how to use the DTCP messages, see “[Example: Using DTCP Messages to Trigger, Verify, and Remove Traffic Mirroring for Subscribers](#)” on page 615.

RELATED DOCUMENTATION

| [Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview](#) | 590

Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring

DTCP-initiated and RADIUS-initiated subscriber secure policy mirroring both use the radius-flow-tap service. If you remove the **radius-flow-tap** configuration, then both types of mirroring are disabled. You

can use the **dtcp-only** statement to cause RADIUS requests to initiate mirroring for a subscriber to be rejected; the mirroring service is not activated. The statement has no affect on DTCP-based mirroring.

Existing RADIUS-initiated mirroring is not affected by the statement, so to be effective you must issue the statement before a RADIUS-initiated service is activated for the subscriber. DTCP-initiated mirroring and FlowTapLite services, which use DTCP, are not affected.

To prevent RADIUS requests from initiating mirroring:

- Enable only DTCP support.

```
[edit system services]
user@host# set dtcp-only
```

RELATED DOCUMENTATION

[Subscriber Secure Policy Overview | 553](#)

[Configuring Support for Subscriber Secure Policy Mirroring | 574](#)

Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs

Example: Configuring Traffic That Is Mirrored Using DTCP-Initiated Subscriber Secure Policy

IN THIS SECTION

- [Requirements | 598](#)
- [Overview | 599](#)
- [Configuration | 599](#)

This example shows how to configure traffic that is mirrored using DTCP-initiated subscriber secure policy.

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 203.0.113.0/8 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 DTCP ADD 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 DTCP ENABLE command. To stop filtering traffic that is currently being mirrored:

- Send a DTCP DELETE message to remove the current policy.
- Modify the configuration with the new version of the policy.
- Send a DTCP ADD message to add the policy.
- Send a DTCP ENABLE message to enable the policy.

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 203.0.113.0/8
```

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 {
          203.0.113.0/8;
        }
        protocol tcp;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

RELATED DOCUMENTATION

[Subscriber Secure Policy Overview | 553](#)

[Configuring Support for Subscriber Secure Policy Mirroring | 574](#)

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 \(DTCP\) | 609](#)

[DTCP Messages Used for Subscriber Secure Policy | 585](#)

Configuring DTCP Messages Used for DTCP-Initiated Subscriber Secure Policy Mirroring

IN THIS CHAPTER

- [ADD \(DTCP\) | 604](#)
- [DELETE \(DTCP\) | 609](#)
- [ENABLE \(DTCP\) | 611](#)
- [LIST \(DTCP\) | 613](#)
- [Example: Using DTCP Messages to Trigger, Verify, and Remove Traffic Mirroring for Subscribers | 615](#)

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
Flags: flag
Seq: sequence-number
Authentication-Info: ssh-authentication-string
```

Description

Specify the DTCP attributes that do one of the following:

- Trigger the router to initiate traffic mirroring.
- 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

This DTCP message is supported for both FlowTapLite and radius-flow-tap services.

NOTE: Starting with Junos OS Release 12.3, DTCP ADD requests are validated for the IP version. The source IP and destination IP addresses must contain a matching IP address family, which must match with the value of the IPVersion field if it is available in the ADD message.

NOTE: Consult the documentation for your mediation device to learn how to configure DTCP messages on the device.

BEST PRACTICE: The Account Session ID, Interface Identifier, and Subscriber User Name trigger attributes are optimized for a scaled subscriber management environment. Forwarding of mirrored traffic begins almost immediately when you include one or more of these three attributes and none of the non-optimized attributes in DTCP ADD messages.

If you include any of the non-optimized trigger attributes in the DTCP ADD message in a scaled subscriber management environment, some delay might be observed between the time when the DTCP ADD message is sent and the time when forwarding starts for the mirrored traffic. For example, if there are 10,000 subscriber sessions on the router, forwarding of the mirrored traffic might be delayed for less than one minute. This delay occurs when you specify any non-optimized attribute, with or without any optimized attribute. The delay occurs regardless of the order of attributes in the DTCP packet.

When a subscriber matches more than one of the DTCP mirroring triggers in an ADD message, the router processes the triggers in the following order:

1. **X-Act-Sess-Id**
2. **X-Call-Sta-Id**
3. **X-IP-Addr**
4. **X-Interface-Id**
5. **X-NAS-Port-Id**
6. **X-RM-Circuit-Id**
7. **X-UserName**

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 **X-RM-Circuit-Id**, 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.

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.

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.

- **X-Act-Sess-Id**—Text string of the accounting session ID associated with the subscriber session. The intercept terminates when the subscriber logs out.

BEST PRACTICE: We recommend that you include other triggers to ensure that all sessions for a subscriber are intercepted.

- **X-Call-Sta-Id**—Text string of the calling station ID associated with the subscriber. If the subscriber is not logged in, the policy is applied at any current or subsequent subscriber log in.
- **X-IP-Addr**—IPv4 address that is associated with the interface for a subscriber.

If the subscriber is not using the default logical system, you must also include the **X-Logical-System** attribute in your DTCP message. If the subscriber is not using the default routing instance, you must also include the **X-Router-Instance** attribute in your DTCP message.

- **X-Interface-Id**—Interface description string on which traffic mirroring is performed. Traffic is mirrored for all subscribers that use this interface; for example, **ge-0/0/0.1** or **demux0.107472834**.
- **X-NAS-Port-Id**—Text string of the NAS port ID associated with the subscriber.
- **X-RM-Circuit-Id**—For PPPoE subscribers, the agent circuit ID (ACI) in the PPPoE Intermediate Agent (PPPoE IA) tag.

For DHCP subscribers, use **X-RM-Circuit-Id** with the agent remote ID (ARI), **X-RM-Agent-Id**, to completely specify a trigger for the DHCP option 82 value that is associated with this session.

- **X-RM-Agent-Id**—For PPPoE subscribers, the agent remote ID (ARI) in the PPPoE IA tag.

For DHCP subscribers, **X-RM-Agent-Id** is the option 82 Agent-Remote-ID suboption and you can use it alone as a trigger. You can also use it with the ACI, **X-RM-Circuit-Id**, to completely specify a trigger for the DHCP option 82 value that is associated with this session.

- **X-Logical-System**—Include in addition to the **X-IP-Addr** or **X-UserName** attribute for subscribers that use anything other than the default logical system. **X-Logical-System** is ignored if neither of those attributes is included in the message. The default logical system is assumed when **X-Logical-System** is not included in the ADD message.
- **X-Router-Instance**—Include in addition to the **X-IP-Addr** or **X-UserName** attribute for subscribers that use anything other than the default routing instance. **X-Router-Instance** is ignored if neither of those attributes is included in the message. The default routing instance is assumed when **X-Router-Instance** is not included in the ADD message.
- **X-UserName**—Subscriber's user name. For subscribers not using the default logical system or routing instance, you can also include the **X-Logical-System** or **X-Router-Instance** attributes.

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 Messages Used for Subscriber Secure Policy | 585](#)

[Packet Header for Mirrored Traffic Sent to Mediation Device | 586](#)

Sample Output

```
ADD DTCP/0.7
Csource-ID: ft-user1
Cdest-ID: cd1
Priority: 2
X-JTap-Cdest-Dest-Address: 203.0.113.50
X-JTap-Cdest-Dest-Port: 7890
X-JTap-Cdest-Source-Address: 203.0.113.9
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
```

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

Remove traffic mirroring for a subscriber. Mirroring of the existing subscriber is stopped. This DTCP message is supported for both FlowTapLite and radius-flow-tap services.

NOTE: Consult the documentation for your mediation device to learn how to configure DTCP messages on the device.

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 remove 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 Messages Used for Subscriber Secure Policy | 585](#)

[DTCP-Initiated Subscriber Secure Policy Overview | 579](#)

List of Sample Output

[DELETE DTCP on page 610](#)

Sample Output

The following sample shows how to remove 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
```

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.

This DTCP message is supported for the radius-flow-tap service. It is not supported for the FlowTapLite service.

NOTE: Consult the documentation for your mediation device to learn how to configure DTCP messages on the 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

[DTCP Messages Used for Subscriber Secure Policy | 585](#)

[DTCP-Initiated Subscriber Secure Policy Overview | 579](#)

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
```

LIST (DTCP)

Syntax

```
LIST DTCP/0.7
Csource-ID: user-name
Cdest-ID: variable
Flags: BOTH
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. This DTCP message is supported for both FlowTapLite and radius-flow-tap services.

NOTE: Consult the documentation for your mediation device to learn how to configure DTCP messages on the device.

NOTE: You can ignore the following fields in command output: AVERAGE-BANDWIDTH, MATCHING-PACKETS, MATCHING-BYTES, and NUM-REFRESH. The LAST-REFRESH field shows the time when the corresponding DTCP ADD request was processed.

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.

If a LIST DTCP command is sent with multiple Cdest-IDs, the error **400 Bad Request** is displayed.

Flags: *flag*—BOTH is the only flag supported. This field must be included in the LIST message for the LIST request to not be rejected until Junos OS Releases 14.1R4 and 14.2R2. Starting with Junos OS Release 14.1R5, 14.2R3, and 15.1R1, the Flags field is not a required parameter in the DTCP LIST message. The LIST request is not rejected if the LIST message does not contain the Flags field. If the DTCP LIST message contains the Flags field, the value of that field is processed. If the LIST message does not contain the Flags field, the CRITERIA field parameter is used for the Flags field.

Starting with Junos OS Release 12.3, when more than one CDest-ID parameter is present in the DTCP LIST or DELETE DTCP commands, the error code 400 (Bad Request) is returned in the response, instead of the error code 431 (Unknown Content Destination).

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 Messages Used for Subscriber Secure Policy | 585](#)

[DTCP-Initiated Subscriber Secure Policy Overview | 579](#)

List of Sample Output

[LIST DTCP on page 614](#)

Sample Output

LIST DTCP

```
LIST DTCP/0.7
Csource-ID: dtcp1
Cdest-ID: cd1
Flags: BOTH
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: 203.0.113.224
FLAGS: BOTH
AVERAGE-BANDWIDTH: 0
MATCHING-PACKETS: 0
```

```

MATCHING-BYTES: 0
NUM-REFRESH: 0
LAST-REFRESH: 2019-06-13 23:45:34.734
X-JTAP-CDEST-DEST-ADDRESS: 192.0.2.168
X-JTAP-CDEST-DEST-PORT: 65535
X-JTAP-CDEST-SOURCE-ADDRESS: 198.51.100.10
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: 203.0.113.224
FLAGS: BOTH
AVERAGE-BANDWIDTH: 0
MATCHING-PACKETS: 0
MATCHING-BYTES: 0
NUM-REFRESH: 0
LAST-REFRESH: 2019-06-13 23:45:48.912
X-JTAP-CDEST-DEST-ADDRESS: 192.0.2.168
X-JTAP-CDEST-DEST-PORT: 65535
X-JTAP-CDEST-SOURCE-ADDRESS: 198.51.100.10
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 Remove 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.
- Remove traffic mirroring on the two subscriber interfaces.
- Verify that traffic mirroring was stopped on the two subscriber interfaces.

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: 192.0.2.168
X-JTap-Cdest-Dest-Port: 65535
X-JTap-Cdest-Source-Address: 198.51.100.10
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: BOTH
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: 192.0.2.168
X-JTap-Cdest-Dest-Port: 65535
```

```

X-JTap-Cdest-Source-Address: 198.51.100.10
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: 203.0.113.112
X-JTap-Cdest-Dest-Port: 7899
X-JTap-Cdest-Source-Address: 192.0.2.9
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
```

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
CSOURCE-ID: dtcp1
CDEST-ID: cd1
CSOURCE-ADDRESS: 203.0.113.224
FLAGS: BOTH
X-JTAP-CDEST-DEST-ADDRESS: 192.0.2.168
X-JTAP-CDEST-DEST-PORT: 65535
X-JTAP-CDEST-SOURCE-ADDRESS: 198.51.100.10
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: 203.0.113.224
FLAGS: BOTH
X-JTAP-CDEST-DEST-ADDRESS: 192.0.2.168
X-JTAP-CDEST-DEST-PORT: 65535
```

```

X-JTAP-CDEST-SOURCE-ADDRESS: 198.51.100.10
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

```

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 removed.

```
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 | 579](#)

[Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview | 590](#)

Configuring Subscriber Secure Policy Support for IPv4 Multicast Traffic

IN THIS CHAPTER

- [Subscriber Secure Policy Support for IPv4 Multicast Traffic | 621](#)
- [Enabling Subscriber Secure Policy Mirroring for IPv4 Multicast Traffic | 622](#)

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 | 622](#)

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 subscriber secure policy to support IPv4 multicast traffic mirroring:

1. Configure the radius-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 | 621](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

[Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview | 590](#)

Configuring Intercept-Related Information for Subscriber Secure Policy

IN THIS CHAPTER

- [Intercept-Related Events Transmitted to the Mediation Device | 623](#)
- [SNMP Traps for Subscriber Secure Policy LAES Compliance | 624](#)
- [Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring | 625](#)
- [Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring | 626](#)

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 with privacy (encryption) configured 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 624](#) describes the supported SNMPv3 traps and their related LAES messages.

RELATED DOCUMENTATION

[Subscriber Secure Policy Overview | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

[Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview | 590](#)

[SNMP Traps for Subscriber Secure Policy LAES Compliance | 624](#)

[Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring | 626](#)

SNMP Traps for Subscriber Secure Policy LAES Compliance

Table 47 on page 624 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 47: Subscriber Secure Policy SNMPv3 Traps for LAES Messages

SNMPv3 Trap	LAES Message	Description
<code>jnxJsPacketMirrorLiSubscriberLoggedIn</code>	<ul style="list-style-type: none"> • 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.
<code>jnxJsPacketMirrorSessionLiSubscriberLogInFailed</code>	<ul style="list-style-type: none"> • 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.
<code>jnxJsPacketMirrorInterfaceLiSubscriberLoggedOut</code>	<ul style="list-style-type: none"> • access-session-end • packet-data-session-end 	A subscriber, who had an active mirrored service, has logged out.
<code>jnxJsPacketMirrorInterfaceLiServiceActivated</code>	<ul style="list-style-type: none"> • packet-data-session-already-established 	A mirrored session has been activated.
<code>jnxJsPacketMirrorSessionLiServiceActivationFailed</code>	–	A mirrored session for a subscriber has failed.
<code>jnxJsPacketMirrorSessionLiServiceDeactivated</code>	–	A mirrored session for an established subscriber has been deactivated.
<code>jnxJsPacketMirrorMirroringFailure</code>	–	<p>A mirrored service request failed due to an invalid value in the request.</p> <p>Note: This trap is not related to LAES messages.</p>

Table 47: Subscriber Secure Policy SNMPv3 Traps for LAES Messages (*continued*)

SNMPv3 Trap	LAES Message	Description
<code>jnxJsPacketMirrorTriggerType</code>	–	The type of trigger that caused the mirroring session to be activated.
<code>jnxJsPacketMirrorCallingStationIdentifier</code>	–	The calling station ID of the subscriber whose traffic is currently being mirrored.
<code>jnxJsPacketMirrorNasIdentifier</code>	–	The NAS ID of the session in which traffic is being mirrored.
<code>jnxJsPacketMirrorTargetIPv6Address</code>	–	The IPv6 address of the subscriber interface that is being mirrored.

RELATED DOCUMENTATION

[Intercept-Related Events Transmitted to the Mediation Device | 623](#)

[Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring | 626](#)

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.

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*.

NOTE: Starting in Junos OS Release 16.1R1, when you configure SNMP trap notifications for subscriber secure policy on MX Series routers, you must configure the target parameters for mediation devices so that the SNMPv3 traps are sent with privacy (encrypted). Targets without privacy configured cannot receive the notifications. In earlier releases, you can configure target parameters without privacy, allowing unencrypted notifications to be sent to the mediation devices.

For more information about configuring subscriber secure policies, see [“Subscriber Secure Policy Overview” on page 553](#).

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 SNMPv3 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 | 623](#)

[SNMP Traps for Subscriber Secure Policy LAES Compliance | 624](#)

[Example: SNMPv3 Traps Configuration for Subscriber Secure Policy Mirroring | 626](#)

[SNMPv3 Overview](#)

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.

```
[edit snmp]
v3 {
  usm {
    local-engine {
      user mediation-device1 { ## Name of the mediation device
        authentication-md5 {
          authentication-key "$ABC123$ABC123"; ## SECRET-DATA
        }
        privacy-des {
          privacy-key "$ABC123"; ## SECRET-DATA
        }
      }
    }
  }
  target-address md1 {
    address 198.51.100.240; ## 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 privacy;
      security-name mediation-device1; ## Name of the mediation device
    }
    notify-filter nf1;
  }
  notify n1 {
    type trap;
    tag mediation-8;
  }
  notify-filter nf1 {
    oid .jnxJsPacketMirrorMIB include;
  }
}
view pkt-mirror-mib oid jnxJsPacketMirrorMIB include
```

Configure the radius-flow-tap service to support subscriber secure policy mirroring.

```
[edit services radius-flow-tap]
```

```
logical-system LS1 routing-instance RI1
snmp {
    notify-targets ip-address;
}
source-ipv4-address 198.51.100.255
```

RELATED DOCUMENTATION

[Subscriber Secure Policy Overview | 553](#)

[Configuring SNMPv3 Traps for Subscriber Secure Policy Mirroring | 625](#)

[SNMPv3 Overview](#)

8

PART

Remote Device and Service Management

[Configuring Remote Device Services Management](#) | **631**

[Configuring TCP Port Forwarding for Remote Subscriber Services](#) | **655**

[Configuring IPFIX Mediation for Remote Device Monitoring](#) | **667**

[Collection and Export of Local Telemetry Data on the IPFIX Mediator](#) | **677**

Configuring Remote Device Services Management

IN THIS CHAPTER

- Remote Device Services Manager (RDSM) Overview | 631
- Configuring Remote Device Management for Service Provisioning | 648
- Reconfiguring a Remote Device for RDSM | 652
- Reloading a Dictionary File for RDSM | 653


Remote Device Services Manager (RDSM) Overview

In some service provider use cases, subscriber services span both a broadband network gateway (BNG) and one or more access nodes. In order to minimize the number of network elements requiring operations support system/business support system (OSS/BSS) integration, the BNG and downstream access nodes are represented to back-office systems as a single, logical system. The service provider's back-office systems provide external configuration and management, authentication, and provisioning for subscriber services on the BNG and its downstream access nodes. To the back-office systems, including PCRF and TACACS+ applications, the BNG and its nodes represent a single addressable network element. The BNG proxies for the downstream devices for service provisioning and deprovisioning.

Starting in Junos OS Release 18.3R1, MX Series routers used as a BNG support remote-device services by means of the remote device services manager (RDSM, using the `rdmd` daemon).

[Figure 23 on page 632](#) shows a sample topology for an MX Series BNG using RDSM. The BNG is connected to OLTs that serve as the downstream, remote devices for provisioning subscriber services, in addition to their conventional role of terminating passive optical network (PON) access per individual subscriber access-lines. The OLTs are logical extensions to the BNG, so that the BNG and its downstream access nodes are presented to back-office systems as a single addressable network element. The BNG uses TCP port forwarding to mediate communications between the remote devices and the back-office system. For more information about TCP port forwarding for remote device management access, see [“TCP Port Forwarding for Remote Device Management” on page 655](#).

Figure 23: Topology for Remote Device Management



The back-office management and provisioning system uses NETCONF XML protocol over SSH for tasks such as base configuration of the remote device before subscriber negotiation begins, configuration of Layer 2 data paths for new subscribers, displaying remote device status, and troubleshooting the remote device. The BNG demultiplexes requests from the management system to the remote devices. Multiple NETCONF sessions can exist to a single remote device.

In this sample topology, the system includes a management platform, PCRF, TACACS+ server, and an IPFIX collector:

- The PCRF sources the subscriber services that are provisioned locally on the MX BNG locally and remotely on the OLTs.
- The TACACS+ server is used to authenticate and validate access to the remote device, perform system accounting, and control operator access. The remote device dynamically initiates a TACACS+ TCP session in response to NETCONF protocol configuration from an external management platform or station. The BNG multiplexes requests from the remote devices to the TACACS+ server.

For remote device access from the back-office system, the server initiates TACACS+ authentication for the following conditions:

- The BNG initiates service configuration for a remote device. The TACACS+ server authenticates the session when the NETCONF TCP socket used by the BNG to provision or deprovision the remote service is opened. After authentication, the session is maintained without authentication or authorization for each remote procedure call (RPC) used for the service action.
- The external management station is used to configure the remote device or access it for monitoring (**show** commands) or troubleshooting.
- The IPFIX collector receives records containing system and connection-level statistics and other information from the MX BNG, which operates as an IPFIX mediator between the OLTs (IPFIX exporters) and the external IPFIX collector. The BNG proxies for the downstream devices. It acts as an IPFIX collector to receive data from the remote devices and as an IPFIX exporter to send data upstream over a single

TCP or TLS session to the collector. For more information about using the BNG as an IPFIX mediator, see [“IPFIX Mediation on the BNG” on page 667](#).

Remote Services

The MX BNG represents a single point of management to external authority for all subscriber services, local and remote. The remote services are also represented by locally configured dynamic service profiles that are referenced by external authority in the same way as local services on the BNG. Consequently, there is a consistent interface between external authority and the BNG for all service actions. The NETCONF XML Management Protocol is used for provisioning and deprovisioning the remote services.

Local subscriber services are defined by dynamic service profiles with zero or more arguments to satisfy subscriber-specific policies. External authorities, such as PCRF, generally use a referential model to provide services. The PCRF charging rule specifies the name of the dynamic service profile and argument values that are applied during subscriber negotiation for service provisioning (activation) or as an update after the subscriber is active. The service is presented to the remote device by the RDSM XML dictionary for that device to parse, interpret, and apply, allowing the charging-rule or service from external authority to be opaquely passed to the remote device with minimal processing. The remote service profile might include one or more variables to define service parameters.

However, remote services can also be applied in a non-referential manner. In this case, the external authority specifies the remote service referentially as it would for a local service. The remote service profile includes one or more variables to define service parameters. The RDSM then uses the data dictionary assigned to the remote device to configure the service on that device. The content of the RDSM dictionary for a device is different depending on whether the service provider uses the referential or non-referential method.

The remote dynamic service profile is very lightweight compared to a local service profile, which can include a large number of configuration stanzas. A remote dynamic service profile contains only two things:

- You must specify that the dynamic profile type is **remote-device-service**. That configuration prevents the profile from being used as a local service profile. This means that you cannot configure a dynamic service profile to be dual-purpose (both local and remote).
- The remote service profile can optionally include a variable stanza to pass argument values to the remote device. The variable stanza can be used for either the referential or non-referential methods.

Any additional configuration fails commit check. Because the remote service profile is so specific, a dedicated service profile is required for each remote service. For the external authority, this means that each remote service requires a separate PCRF charging rule.


Process Flows for RDSM Provisioning and Deprovisioning

Subscriber services are provisioned and deprovisioned as follows:

- Provisioned during subscriber login. The services can be sourced from the PCRF in response to initiation with Gx CCR-I/CCR-A message exchanges.
- Deprovisioned during subscriber logout.
- Provisioned or deprovisioned for active subscriber sessions in response to external authority, such as Gx RAR messages from the PCRF.

Figure 24 on page 634 shows the process flow when RDSM successfully provisions services on three eligible remote devices, OLT1, OLT2, and OLT3, by instantiating the upstreamBW service profile.

Figure 24: RDSM Service Provisioning on a Remote Device: Successful Subscriber Negotiation Flow



1. Service provisioning begins when a subscriber logs in and authd sends a request to RDSM to instantiate the remote service profile on eligible remote devices during the negotiation.
2. RDSM establishes a list of remote devices that are eligible for the service to be provisioned:
 - The Layer 2 access domain for the device must match the subscriber location. The access domain consists of a configured list of VLAN ranges or individual VLAN IDs. The subscriber's outer VLAN tag must be on this list.
 - The NETCONF TCP connection to the device must be up. Although a device in the down state is not eligible for provisioning, it might be available for reconfiguration if it transitions later to the up state.

3. RDSM performs an initial validation before it responds to the remote service profile instantiation request:

- When validation passes, RDSM sends a service profile instantiation pending ACK response to authd. The service provisioning is now pending.
- If validation fails, RDSM returns a NACK response to authd and abandons service provisioning.

The validation checks performed by RDSM typically do not fail for active subscriber sessions. Reasons for failure include the following:

- No remote device has a subscriber location that matches the access domain.
- The dictionary located on the BNG does not include an entry for the requested remote service profile. Consequently there are no RPCs to provide the service variables and install the service.

4. RDSM resolves any required parameters for each remote device; at a minimum, this includes the subscriber identifier.

5. RDSM then uses the dedicated NETCONF session to each of the eligible devices to issue a series of RPC calls as specified in the dictionary for provisioning the service.

Service provisioning takes place in parallel for the eligible devices. Provisioning fails for a device when either of the following occurs:

- The RDSM receives an explicit error for any RPC call.
- The response times out.

The following ERRMSG event is logged in either case:

remote device *device-name* ip-address service *service-name* provisioning failed for subscriber *subscriber-id*

6. Remote devices that are successfully provisioned return an ACK response to RDSM.


If one or more remote devices fails to be provisioned, RDSM rolls back the service on every remote device that was successfully provisioned. RDSM uses the dedicated NETCONF session to each of these devices to issue a series of RPC calls as specified in the dictionary for deprovisioning the service.

7. RDSM sends an out-of-band notification to authd to report whether the remote service was provisioned on the remote devices.

- When provisioning is successful for all remote devices, RDSM sends a service provisioning complete response to authd.
- If one or more of the eligible remote devices fails to be provisioned, RDSM reports a provisioning failure to authd.

Figure 25 on page 636 shows the process flow when RDSM successfully updates subscriber services on three eligible remote devices, OLT1, OLT2, and OLT3 by instantiating the upstreamBW service profile with different parameter values than were used during login.

Figure 25: RDSM Service Provisioning on a Remote Device: Subscriber Update Flow



Updating subscriber services begins when authd sends a request to RDSM to instantiate the remote service profile to update the service. The process flow is the same as for the subscriber login flow, except that RDSM does not respond to the instantiation request until all processing required to provision the service is complete. That means that when the validation check passes, RDSM does not send a service profile instantiation pending ACK response to authd; if validation fails, RDSM does return a NACK response to authd and abandons service deprovisioning.

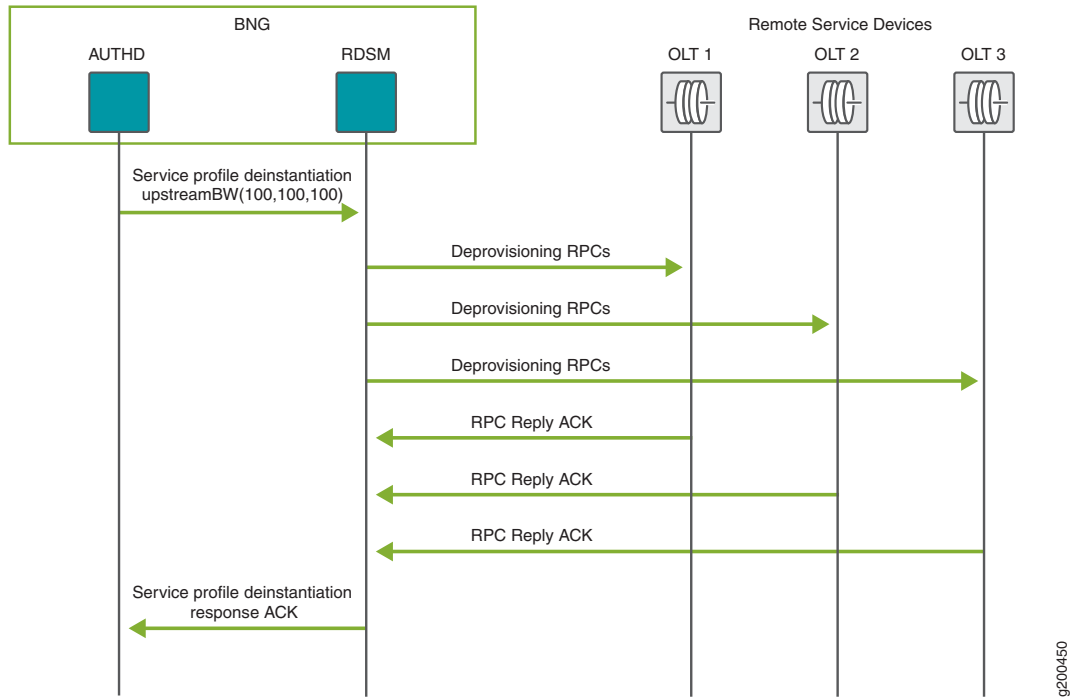
[Figure 26 on page 637](#) shows the process flow when RDSM successfully deprovisions services to update the three eligible remote devices, OLT1, OLT2, and OLT3, by deinstantiating the upstreamBW service profile.

The deprovisioning process flow is the same for both a subscriber logout and an update request from authd.

RDSM does not respond to authd until all required processing to deprovision the service has completed (including any retry of failures); this allows subscriber logout to proceed regardless of the deprovisioning outcome.

Service deprovisioning typically does not fail; if it does, then you may have to take some corrective action on the remote device for deprovisioning to succeed.

Figure 26: RDSM Service Deprovisioning on a Remote Device: Subscriber Logout and Update Flow



1. Service deprovisioning begins when either of the following occurs:
 - A subscriber logs out and authd sends a request to RDSM to deinstantiate the remote service profile on eligible remote devices.
 - authd sends an update request to RDSM to deinstantiate the remote service profile on eligible remote devices.
2. RDSM maintains a list of remote devices that are provisioned with the service. If the NETCONF TCP connection to the device is down, deprovisioning is not attempted because it is assumed to have occurred by some other means. For example, the device may have been reconfigured with a default, baseline configuration and subsequent operator action initiated reconfiguration by the BNG for all active subscriber services.
3. RDSM performs an initial validation before it responds to the remote service profile deinstantiation request:
 - When validation passes, RDSM does not send a response to authd.
 - If validation fails, RDSM returns a NACK response to authd and abandons service deprovisioning. Reasons for a validation failure include the following:
 - No configured remote device is in the up state.

- The dictionary located on the BNG does not include an entry for the requested remote service profile or deprovisioning action. Consequently there are no RPCs to provide the service variables and remove the service.

The validation checks performed by RDSM typically do not fail for active subscriber sessions.

4. RDSM resolves any required parameters for each remote device; at a minimum, this includes the subscriber identifier.
5. RDSM then uses the dedicated NETCONF session to each of the eligible devices to issue a series of RPC calls as specified in the dictionary for deprovisioning the service. Service deprovisioning takes place in parallel for the eligible devices. Deprovisioning fails for a device when either of the following occurs:
 - The RDSM receives an explicit error for any RPC call.
 - The response times out.

In either case, RDSM retries the deprovisioning action up to 5 times, at 5-second intervals. If the attempts all fail, then the following ERRMSG event is logged in either case:

remote device *device-name* ip-address service *service-name* de-provisioning failed for subscriber *subscriber-id*

6. Remote devices that are successfully deprovisioned return an ACK response to RDSM.
7. RDSM sends an out-of-band notification to authd to report whether the remote service was deprovisioned on the remote devices.
 - When deprovisioning is successful, RDSM sends a service deprovisioning complete response to authd, which then completes the subscriber logout.

In the case of an update request rather than a subscriber logout, RDSM sends a service profile deinstantiation complete response to authd, which completes the service session clean-up.

 - If one or more of the eligible remote devices fails to be deprovisioned, RDSM reports a deprovisioning failure to authd.

RDSM Dictionary for Implementing Service Actions

An XML dictionary stored locally on the BNG is an integral component of remote device service management. Each remote service provisioned by the external authority must have an entry in an RDSM dictionary on the BNG. The dictionary translates the PCRF-sourced charging rule to a set of vendor-specific remote procedure calls (RPCs) in the entry associated with the service. The RPCs then provision or deprovision the service. Because the RPCs are vendor-specific, so is the dictionary. This means that separate

dictionaries are required for each vendor's remote device. For a given vendor's devices, different software releases on the devices may require different dictionaries as well.

The dictionary format is sufficiently flexible to support both referential services and non-referential services, where:

- A referential service means that the entire service, including arguments, is presented opaquely to the remote device as received from external authority via the RDSM dictionary. The dynamic service profile can include a variable stanza that is used by the dictionary during translation of the arguments. The remote device parses, interprets and applies the arguments on its own without any interpretation or parsing by the BNG.
- A non-referential service means that all arguments supplied by the external authority must be resolved and provided to the remote device individually by one or more RPCs. In this case, the dynamic service profile may require a variable stanza that is used by the dictionary during translation of the arguments.

In either case, the dictionary must specify the means—typically a Layer 2 location—to identify the subscriber suitable for the remote device to distinguish one subscriber from another.

The XML RDSM dictionary has the following general format:

```
<junos-rdm-dictionary>
  <junos-rdm-parameters>
    <junos-rdm-parameter>
      <junos-rdm-name>...</junos-rdm-name>
      <junos-rdm-source>...</junos-rdm-source>
      <junos-rdm-index>...</junos-rdm-index>
    </junos-rdm-parameter>

    <junos-rdm-parameter>
...
  </junos-rdm-parameter>
</junos-rdm-parameters>

  <junos-rdm-services>      <junos-rdm-service>
    <junos-rdm-name>...</junos-rdm-name>
    <junos-rdm-provision>
      <junos-rdm-service-configuration>
...
      </junos-rdm-service-configuration>
    </junos-rdm-provision>
    <junos-rdm-deprovision>
      <junos-rdm-service-configuration>
...
      </junos-rdm-service-configuration>
    </junos-rdm-deprovision>
```

```

    </junos-rdm-service>
</junos-rdm-services>

<junos-rdm-open-configuration>
  <junos-rdm-rpc>...</junos-rdm-rpc>
  <junos-rdm-rpc>...</junos-rdm-rpc>
</junos-rdm-open-configuration>

<junos-rdm-edit-configuration>
  <junos-rdm-rpc>...</junos-rdm-rpc>
  <junos-rdm-rpc>...</junos-rdm-rpc>
  <junos-rdm-rpc>...</junos-rdm-rpc>
</junos-rdm-edit-configuration>

<junos-rdm-commit-configuration>
  <junos-rdm-rpc>...</junos-rdm-rpc>
  <junos-rdm-rpc>...</junos-rdm-rpc>
</junos-rdm-commit-configuration>

<junos-rdm-close-configuration>
  <junos-rdm-rpc>...</junos-rdm-rpc>
  <junos-rdm-rpc>...</junos-rdm-rpc>
</junos-rdm-close-configuration>
</junos-rdm-dictionary>

```

Table 48 on page 640 defines the individual components of the dictionary.

Table 48: Definitions of XML Dictionary Components

junos-rdm-parameters	Parameter block that lists individual parameters that configure the service.
junos-rdm-parameter	Individual parameter.
junos-rdm-name	In the parameter block, this element identifies the subscriber on the remote device or the PCRF argument. Use the subscription-id for the subscriber and the name of the argument for any argument specified in the PCRF.
junos-rdm-source	Source of the parameter value: <ul style="list-style-type: none"> • subscriber-session when the value is sourced from the SDB session. • service-profile when the value is sourced from the service profile argument.

Table 48: Definitions of XML Dictionary Components (*continued*)

junos-rdm-index	<p>Index, such as an enumerated type value, that resolves the parameter from the specified source. The subscriber-session source requires this to map the parameter to an SDB attribute used to resolve the parameter value.</p> <p>For example, for some use cases, the PCRF subscription-id is stored in the subscriber SDB entry that is referenced by an index (attribute type) to resolve this parameter.</p>
junos-rdm-services	Service block that lists one or more remote services supported by the device.
junos-rdm-service	Individual remote service defined by service name, provisioning configuration, and deprovisioning configuration.
junos-rdm-name	In the service block, this element is the name of the service. It is the base service name, without arguments, of the service sourced from the PCRF.
junos-rdm-provision	Provisioning block that includes provisioning configuration.
junos-rdm-deprovision	Deprovisioning block that includes deprovisioning configuration.
junos-rdm-service-configuration	Service configuration that includes one or more RPCs to provision or deprovision the service. When arguments are specified in the PCRF service for provisioning, the RPCs include those arguments.
junos-rdm-open-configuration	Block that includes zero or more RPCs to begin configuration of the remote device.
junos-rdm-edit-configuration	Block that includes one or more RPCs to edit the configuration and apply service provisioning or deprovisioning actions to the device in bulk, by referencing the junos-rdm-provision or junos-rdm-deprovision block for the specified service. The configuration for each service that is part of the bulk update to the remote device is included.
junos-rdm-commit-configuration	Block that includes zero or more RPCs to commit the edits to the remote device.

Table 48: Definitions of XML Dictionary Components (*continued*)

junos-rdm-close-configuration	Block that includes zero or more RPCs to end configuration of the remote device.
junos-rdm-rpc	Individual RPC to configure the remote device.

For remote device configuration, the edit configuration is always required to provision or deprovision the service. In some use cases, the open, commit, and close configuration blocks might be optional.

Additional Features for Use with an RDSM Access Model

The features in this section are not required for RDSM, but may be useful in certain use cases or topologies.

A locally generated username is used in interactions with an external authority to authenticate dynamic VLAN, DHCPv4, and DHCPv6 subscribers. Typically, subscriber VLAN tags are included in the username by configuring the **interface-name** option for the **username-include** statement.

Similarly, subscriber VLAN tags are included in the subscription identifier for PCRF interactions by configuring the **interface-name** option for the **subscription-id-data-include** statement.

By convention, the interface name has the following format in both cases:

```
underlying-IFD-name:outer-vlan-tag[-inner-vlan-tag]
```

For some use cases with the RDSM access model, the outer VLAN tag is unique across the system. This means that you can use a different format that excludes the underlying IFD name:

```
outer-vlan-tag[-inner-vlan-tag]
```

To generate the username format without the underlying IFD name, you specify the **vlan-tags** option instead of the **interface-name** option with the **username-include** statement. See *Configuring VLAN Interface Username Information for AAA Authentication* and *Creating Unique Usernames for DHCP Clients* for more information.

To generate the subscription ID format without the underlying IFD name, you specify the **vlan-tags** option instead of the **interface-name** option with the **subscription-id-data-include** statement. See *Configuring the PCRF Partition* for more information.

Some customer networks might have more than one deployment model or use case that results in the MX Series BNG for each case interacting with the same PCRF back-end. In this situation, you might need to distinguish between the use cases for the PCRF.

The Diameter Capability Exchange messages between peers carry the Diameter Product-Name AVP. You can configure nondefault values for the use cases so the PCRF can discriminate between the messages.

See *Messages Used by Diameter Applications and Diameter AVPs and Diameter Applications* for more information.

Response to the External Authority by authd on Success or Failure

How authd responds to the external authority depends on the following:

- The operation being performed for example, provisioning during subscriber login versus updating an existing subscriber session.
- The external authority, Gx (PCRF).

[Table 49 on page 643](#) describes how the authd response varies when the service provisioning or deprovisioning actions are successful.

Table 49: How authd Responds to External Authority When Service Actions Succeed

Operation	Gx
Login	<p>authd initiates CCR-I/CCA-I message exchange to provision the subscriber session.</p> <p>When all services in the CCA-I are provisioned, authd sends a CCR-U message that indicates the service is active for each charging-rule in the CCA-I. Status reporting for local dynamic services is delayed until remote services provisioning completes.</p>
Update	<p>Deprovisioning is applied before provisioning for services included in the same PCRF RAR message.</p> <p>When deprovisioning and provisioning is completed for all service actions included in the PCRF RAA, authd sends an RAA response with a Rule-Report that indicates the service inactive/active state for each charging rule specified in the RAR.</p> <p>Status reporting for local dynamic services is delayed until remote services processing completes.</p>
Logout	<p>authd initiates a CCR-T/CCA-T message exchange to notify PCRF of subscriber termination.</p> <p>authd initiates deprovisioning for all services configured for the subscriber session.</p>
Service device in the up state after the reconfigure command is issued	<p>authd takes no further action when it receives an out-of-band notification from RDSM that the service action succeeded.</p> <p>For example, it does not send a CCR-U message that indicates the service is active for the corresponding charging rule.</p>

[Table 50 on page 644](#) describes how the authd response varies when the service provisioning or deprovisioning actions fail.

Table 50: How authd Responds to External Authority When Service Actions Fail

Operation	Gx
Login	<p>authd initiates a CCR-I/CCA-I exchange to provision the subscriber session.</p> <p>When authd receives notification of failure in the service profile instantiation response or out-of-band from RDSM, authd stops processing any remaining services.</p> <p>authd sends a CCR-U message that reports the following:</p> <ul style="list-style-type: none"> • Service is active for each charging-rule in the CCA-I that successfully provisioned • Service is inactive for the charging rule in the CCA-I that failed provisioning. • Service is inactive for all charging rules not processed because of the failure. <p>authd allows the subscriber session negotiation to complete and reach the active state.</p>
Update	<p>Deprovisioning is applied before provisioning for services included in the same PCRF RAR message.</p> <p>The process varies depending on the actions that fail.</p> <ul style="list-style-type: none"> • When authd receives notification of failure in the service profile deinstantiation response, meaning that RDSM has performed all retries without success, authd continues to process the next service action. <p>This means that when only service deprovisioning fails, the update proceeds and completes.</p> <ul style="list-style-type: none"> • When authd receives notification of failure in the service profile instantiation response, authd stops processing any remaining services. <p>All provisioned and deprovisioned services in the request are rolled-back. That means that services that were successfully provisioned are now deprovisioned. Services that were successfully deprovisioned are now reprovisioned.</p> <p>When all rollback actions are completed, authd sends an RAA response with a Rule-Report that indicates the service inactive/active state for each charging rule specified in the RAR.</p> <p>This means that reprovisioned charging-rules are reported as active and deprovisioned charging-rules are reported as inactive.</p>
Logout	<p>authd initiates a CCR-T/CCA-T message exchange to notify PCRF of subscriber termination.</p> <p>authd initiates deprovisioning for all services configured for the subscriber session.</p> <p>When authd receives notification of failure in the service profile instantiation response or out-of-band from RDSM, authd continues with the logout, including deprovisioning any remaining services.</p>

Table 50: How authd Responds to External Authority When Service Actions Fail (*continued*)

Operation	Gx
Last service device in down state after the reconfigure command is issued	<p>authd takes no further action when it receives an out-of-band notification from RDSM that the service action failed.</p> <p>For example, it does not send a CCR-U that indicates the service is inactive for the corresponding charging rule.</p> <p>Affected subscriber sessions are maintained.</p>

Operator Reconfiguration of Remote Devices

In some circumstances, you might need to manually provision services on a remote device to resynchronize the device with all matching subscriber services that are active and configured on at least one other remote device. Manual provisioning is required in the following scenarios:

- A new remote device is connected to the BNG after one or more subscriber sessions have been negotiated on other remote devices and remote services have been provisioned on those devices.
- The NETCONF session to a remote device with one or more provisioned remote services transitions to the down state, then later recovers and transitions back to the up state. This is effectively the same as a new device being connected to the BNG.

After the NETCONF session is established to the remote device in either of these situations, an ERRMSG event is logged that the device is up. No remote services are currently provisioned on the device. RDSM establishes a list of subscriber remote services that are eligible to be provisioned on the device. These services must be either active or in the process of being provisioned. A separate ERRMSG event is logged, indicating that services are pending reconfiguration:

remote device *device-name ip-address* has *number* services pending reconfiguration

You use the **request services remote-device-management reconfigure service-device** command to provision all active (or in process) subscriber services that map to the access domain associated with the device. The reconfiguration request triggers bulk provisioning of services on the device. If the provisioning of one service fails, the entire bulk provisioning is considered a failure and any successfully provisioned services are rolled back. In this case you have to issue the command again. The rollback applies only to each bulk provisioning attempt, so you can control the effects of a bulk provisioning failure by setting a bulk limit.

NOTE: The remote device is eligible to be automatically provisioned with subscriber services without operator intervention for subscriber logins that occur after the NETCONF session is established.

You can issue a reconfiguration request at any time when the remote device is up. When remote device reconfiguration begins, any new service actions resulting from new subscriber negotiation or existing subscriber update or logout are delayed for the remote device until reconfiguration completes. Also, a reconfiguration request may be performed at any time when the remote device is up. This means that a remote device may be connected to the network and accept new subscriber services provisioning before existing subscribers are provisioned by the reconfiguration request.

The following steps show the RDSM process flow for reconfiguration requests:

1. RDSM maintains a list of remote devices that are provisioned with the service. If the NETCONF TCP connection to the device is down, deprovisioning is not attempted because it is assumed to have occurred by some other means. For example, the device may have been reconfigured with a default, baseline configuration and subsequent operator action initiated reconfiguration by the BNG for all active subscriber services.
2. RDSM performs the following as a bulk operation, where the bulk size maybe up to total number of subscriber services to be provisioned:
 - a. Validates the service before it responds to the reconfiguration request. For example, validation fails when the dictionary located on the BNG does not include an entry for the requested remote service profile or provisioning action, because there are no RPCs to provide the service variables and add the service.
 - b. Resolves any required parameters for each remote device; at a minimum, this includes the subscriber identifier.
 - c. Uses the dedicated NETCONF session to the remote device to issue a series of RPC calls as specified in the dictionary for provisioning the service. Provisioning fails for a device when either of the following occurs:
 - The RDSM receives an explicit error for any RPC call.
 - The response times out.

In either case, RDSM rolls back all service that were successfully provisioned by the bulk operation, reconfiguration is abandoned and RDSM logs the following ERRMSG event:

remote device *device-name* *ip-address* reconfiguration failed

3. If provisioning completes for all the subscriber services on the remote device, RDSM logs the following ERRMSG event:

remote device *device-name* *ip-address* reconfiguration succeeded

External Notification for Service Processing ERRMSG Events

Table 51 on page 647 lists the ERRMSG events that authd can communicate to external management systems and the information that is included in the notifications. Successful remote service actions are only reported to an external authority and do not generate an ERRMSG log.

Table 51: Information Included in External Notifications for ERRMSG Events

ERRMSG Event	Device Name	IP Address	Current State	Number of Services Pending Reconfiguration	Service Name	Subscriber Identifier
Remote device status change from up to down or down to up	✓	✓	✓	–	–	–
Remote device has services pending reconfiguration	✓	✓	–	✓	–	–
Remote device reconfiguration completion (success or failure)	✓	✓	✓	–	–	–
Subscriber remote service provisioning failure	✓	✓	–	–	✓	✓
Subscriber remote service deprovisioning failure	✓	✓	–	–	✓	✓

Benefits of Remote Device Service Management

- Enables topologies where subscriber services span both the MX Series BNG and its access nodes to form a single, logical system.
- Simplifies BNG and remote device configuration and management in topologies that use external management and provisioning systems. The remote devices typically have private addresses unknown to the external system, so the external system addresses only the MX Series BNG.
- Adds a new service profile type for remote services to easily differentiate remote and local services.

RELATED DOCUMENTATION

[Configuring Remote Device Management for Service Provisioning | 648](#)

[Reconfiguring a Remote Device for RDSM | 652](#)

Configuring Remote Device Management for Service Provisioning

You must configure both dynamic service profiles and remote devices. A dynamic service profile is identified for RDSM by configuring the profile type as `remote-device-service`. This profile type prevents the profile from being applied locally on the router. It is limited to application on an external device by RDSM. The external authority, such as PCRF can reference this profile to provision or deprovision services on the remote device.

The remote device configuration includes the device IP address and the dictionary path. The remote device must have an entry in an XML dictionary hosted on the MX BNG. The dictionary translates the service action instructions from the external authority to a set of vendor-specific remote procedure calls (RPCs) in the entry associated with the service. The RPCs then provision or deprovision the service.

Finally, you can configure several parameters for the provisioning method, the NETCONF XML protocol. You must configure the username and password used to access the remote device. Other parameters are optional.

NOTE: Although the following procedure shows only configuration at the `[edit system services]` hierarchy level, and therefore the default routing instance, you can also configure RDSM at the `[[edit routing-instances routing-instance-name system services]` hierarchy level.

You must also configure the back-office system that provides the external authority and management platform for remote device service management. That configuration is outside the scope of this topic. Consult the vendor documents for your back-office equipment.

To configure remote device service management:

1. Configure one or more dynamic service profiles. Specify that the dynamic service profile containing this statement is not applied locally to the router. Instead, it is applied to an external device by means of the remote device services manager daemon (`rdmd`). It enables an external authority, PCRF to reference the dynamic service profile to provision or deprovision services (charging rules) on the remote device.

```
[edit dynamic-profiles profile-name]  
user@host# set profile-type remote-device-service
```

2. Configure one or more devices for remote services.

```
[edit system services remote-device-management]
user@host# edit service-device device-name
```

3. (Optional) Configure the Layer 2 access domain for the remote device.

```
[edit system services remote-device-management service-device device-name]
user@host# set access-domain vlan-id-list [vlan-id-low-vlan-id-high vlan-id]
```

4. Configure the address for the remote device.

```
[edit system services remote-device-management service-device device-name]
user@host# set address ip-address
```

5. Specify the absolute file path for the XML dictionary.

```
[edit system services remote-device-management service-device device-name]
user@host# set dictionary absolute file path
```

6. Specify the provisioning method (only NETCONF XML protocol is supported).

```
[edit system services remote-device-management service-device device-name]
user@host# edit provisioning-method netconf
```

7. Configure the provisioning parameters for the NETCONF protocol. You must specify the username and password; all other options have default values.
 - a. Specify the name used to access the remote device during service management. The maximum length of the name is 64 bytes.

NOTE: If you change the username when any active subscriber services are mapped to the device, the change takes effect only when the device reconnects.

```
[edit system services remote-device-management service-device device-name provisioning-method
netconf]
user@host# set user-name name
```

- b. Specify the password used by the NETCONF protocol to access the remote device during service management. The maximum length of the password is 64 bytes.

NOTE: If you change the password when any active subscriber services are mapped to the device, the change takes effect only when the device reconnects.

```
[edit system services remote-device-management service-device device-name provisioning-method
netconf]
user@host# set password password
```

- c. (Optional) Specify the period during which multiple services are provisioned or deprovisioned based on the assigned dictionary before the configuration is committed to the service device. When the interval times out, the service actions are committed in bulk before additional actions for the device can take place.

NOTE: You can use the **bulk-interval** and **bulk-limit** options together to optimize your service device configuration during scaled subscriber negotiation and service provisioning or subscriber termination and service deprovisioning.

```
[edit system services remote-device-management service-device device-name provisioning-method
netconf]
user@host# set bulk-interval milliseconds
```

- d. (Optional) Specify how many services can be provisioned or deprovisioned during the bulk interval before the configuration is committed to the device.

```
[edit system services remote-device-management service-device device-name provisioning-method
netconf]
user@host# set bulk-limit number
```

- e. (Optional) Specify how long RDSM waits between successive attempts to establish a NETCONF session with the remote device.

```
[edit system services remote-device-management service-device device-name provisioning-method
netconf]
```

```
user@host# set connection-retry-interval seconds
```

- f. (Optional) Specify the TCP port number for the NETCONF session with the remote device.

```
[edit system services remote-device-management service-device device-name provisioning-method
netconf]
user@host# set port port-number
```

- g. (Optional) Specify how many services can be provisioned or deprovisioned as a result of a reconfiguration before the configuration is committed to the service device. When the limit is reached, the service actions are committed in bulk before additional actions for the device can take place.

```
[edit system services remote-device-management service-device device-name provisioning-method
netconf]
user@host# set reconfigure-bulk-limit number
```

- h. (Optional) Specify the period during which the device must respond to an attempt to provision or deprovision a service. The timeout is a failure equivalent to an explicit failure response received from the device.

```
[edit system services remote-device-management service-device device-name provisioning-method
netconf]
user@host# set response-timeout seconds
```

- i. (Optional) Specify how many consecutive response timeouts can occur before the BNG takes action. The default action is to close and reopen the NETCONF connection.

```
[edit system services remote-device-management service-device device-name provisioning-method
netconf]
user@host# set response-timeout-count number
```

Table 52 on page 652 lists commands you can use to view information about your RDSM configuration and operation.

Table 52: show Commands for Remote Device Services Management

Command	Description
show remote-device-management service-devices	Display information about all remote service devices or a specific remote service device.
show remote-device-management services	Display information about all service sessions or a specific service session on remote service devices.
show remote-device-management statistics	Display a global summary of service statistics for all remote devices or detailed statistics for a specific remote service device.
show remote-device-management subscribers	Display information about service sessions for all subscriber sessions or about all service sessions for a specific subscriber session on remote service devices.
show remote-device-management summary	Display summary information about the remote service devices, such as session state and service state.

You can use the [clear remote-device-management statistics](#) command to clear service statistics for all remote devices globally or statistics for a specific remote service device.

RELATED DOCUMENTATION

[Remote Device Services Manager \(RDSM\) Overview | 631](#)

[Reconfiguring a Remote Device for RDSM | 652](#)

[Reloading a Dictionary File for RDSM | 653](#)

Reconfiguring a Remote Device for RDSM

In some circumstances you might need to reconfigure a remote device to manually provision all active subscriber services matching the access domain (list of VLAN ranges and IDs) to which this remote device belongs. The reconfiguration resynchronizes the device with all active (or in process) subscriber services that map to the access domain associated with the device.

For example, if a new remote device is connected to the BNG after subscriber sessions have been brought up on other remote devices in the same access domain and remote services have been provisioned on the devices. The new device is not provisioned at this point, and you would like it be provisioned as if it had been connected during the original service provisioning.

Another situation occurs when the NETCONF session to a provisioned remote device transitions to the down state and then back to the up state. From the perspective of the BNG, this is the same as if the device is new and connected to the BNG for the first time.

You can issue a reconfiguration request at any time when the remote device is up. Reconfiguration provisioning of services occurs in bulk. If the provisioning of one service fails, the entire bulk provisioning is considered a failure and any successfully provisioned services are rolled back. You must issue the command again.

To reconfigure service provisioning for a device:

- Specify the device to be reconfigured.

```
user@host> request services remote-device-management reconfigure service-device device-name
```

The command indicates whether the action succeeds or fails.

RELATED DOCUMENTATION

[Configuring Remote Device Management for Service Provisioning | 648](#)

[Remote Device Services Manager \(RDSM\) Overview | 631](#)

Reloading a Dictionary File for RDSM

You can reload the vendor-specific dictionary to the RDSM database on the BNG by specifying the absolute file path. An example absolute path is `/var/home/dict/remote-device.xml`. The path must end with the `.xml` extension and not exceed 127 characters.

The dictionary defines the set of NETCONF XML protocol commands required to provision, deprovision, and roll back a subscriber service for a remote device. The reload affects all remote service devices that are configured with this dictionary. When you modify an existing dictionary, this is how you apply the updated file.

To reload a dictionary:

- Specify the path for the dictionary to be reloaded.

```
user@host> request services remote-device-management reload-dictionary absolute file path
```

The command indicates whether the action succeeds or fails. A typical cause for failure is when there is an active remote device configured with that dictionary and the device has an active subscriber service.

RELATED DOCUMENTATION

[Remote Device Services Manager \(RDSM\) Overview](#) | 631

[Configuring Remote Device Management for Service Provisioning](#) | 648

Configuring TCP Port Forwarding for Remote Subscriber Services

IN THIS CHAPTER

- [TCP Port Forwarding for Remote Device Management | 655](#)
- [Configure TCP Port Forwarding for Remote Device Management | 659](#)
- [Tracing TCP Port Forwarding Events for Troubleshooting | 662](#)

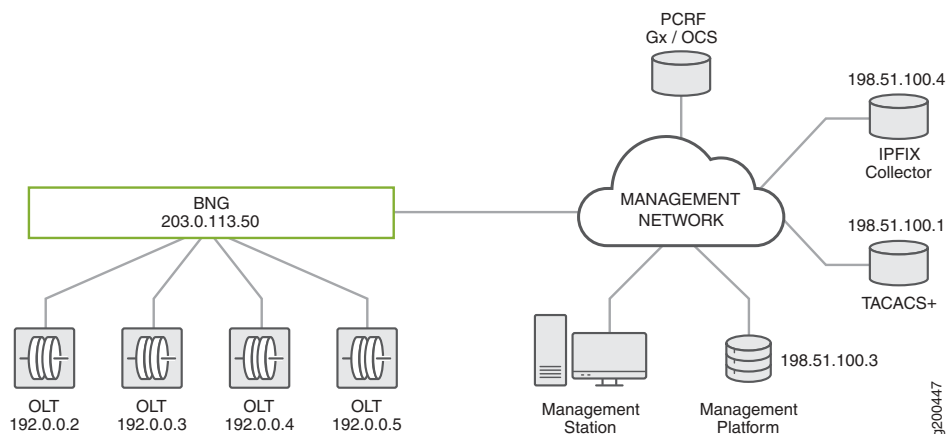
TCP Port Forwarding for Remote Device Management

Port forwarding is a method that enables a router to make a computer or other network device that is connected to it accessible to other computers and network devices from outside of the local network. Port forwarding uses a combination of an IP address and a port number to route network requests to specific devices. This technique is often used to make services on a host or gateway, residing on an internal network, accessible to a host on an external network by remapping the destination IP address and port number for the communication request.

Starting in Junos OS Release 18.3R1, TCP port forwarding (also referred to as TCP forwarding) enables a BNG to mediate communication between its connected access nodes and service provider back-office systems, such as external management and provisioning systems and TACACS+ servers. The BNG and its downstream access nodes are presented to back-office systems as a single addressable network element. You configure unique combinations of listening ports and addresses on the BNG. TCP connections are triggered when traffic from acceptable prefixes arrives on the listening port and matching listening address. Communication requests to and from access nodes are redirected from one address and port number combination to another when packets traverse the MX series router.

Back-office systems use NETCONF XML management protocol over SSH and TACACS+ to exchange requests with access nodes. For provisioning, they can use PCRF and RADIUS to supply service configurations for subscribers. [Figure 27 on page 656](#) shows a sample topology for an external management system use case with optical line terminals (OLTs) connected to the BNG. Similar topologies might have different access nodes, such as DSLAMs, rather than OLTs.

Figure 27: Topology for Remote Device Management



The access nodes in this kind of topology act as logical extensions (remote devices) of the BNG so that the BNG can proxy all external management interactions for them. The BNG is configured with a public address and acts as the single point of management for itself and the access nodes. The remote devices have private addresses and are not publicly accessible. This means that the external systems cannot interact directly with the access nodes. The BNG must be able to mediate management requests between the access nodes and the management system, but it does not need to parse or act on the full content of the requests. This need is met with TCP port forwarding as follows for this use case:

- The external management system uses NETCONF XML protocol over SSH for tasks such as base configuration of the remote device before subscriber negotiation begins, configuration of Layer 2 data paths for new subscribers, displaying remote device status, and troubleshooting the remote device.

In this case, the BNG demultiplexes requests from the management system to the remote devices.

- TACACS+ is used to authenticate and validate access to the remote device, perform system accounting, and control operator access.

In this case, the BNG multiplexes request from the remote devices to the TACACS+ server that works with the external management system.

TCP port forwarding maps one or more combinations of an IPv4 listening address and a TCP port to destination addresses and ports so that the BNG can forward messages appropriately for both use cases. Each mapping is referred to as a *TCP connection pair*. TCP port forwarding operates as follows:

1. When the mapping is configured, the TCP port forwarding process opens the configured listening port and waits for an external system or access node to trigger a connection; that system or node can then be referred to as the *triggering entity*.
2. After the connection between the triggering entity and the BNG is established, TCP port forwarding attempts to open a TCP connection to the other half of the connection pair, which is the forwarding address and port combination defined in the mapping. TCP port forwarding examines only the TCP header information in the management traffic.

3. When both TCP connections have been established, TCP port forwarding monitors the connections for data traffic. When data is received on one connection, it is transmitted on the paired connection.

NOTE:

- If one side of the connection pair closes for any reason, TCP port forwarding closes the paired connection. This connection pair is not reestablished unless the triggering entity makes the connection on the TCP listening port again.
- If a configuration change is made to a TCP mapping while associated connection pairs are active, these connections are closed down. The connections are not reestablished unless the triggering entity makes the connection on the TCP listening port again

TCP port forwarding allows multiple simultaneous TCP connections for any single TCP mapping. You can place a limit on the maximum number of allowed connections.

You can use the following operational commands to manage and monitor TCP port forwarding:

- **clear tcp-forwarding connections**—Enables you to administratively close any current TCP connection pair.
- **clear tcp-forwarding statistics**—Enables you to clear (zero) statistics for the configured TCP mappings and any current TCP connection pairs. You can limit statistics clearing to all connections associated with a specific listening port/listening address combination or to only a single connection pair represented by a specific source address/source port combination. For either combination, you can optionally specify a routing instance; otherwise, the default routing instance is assumed.
- **show tcp-forwarding status**—Displays the status of TCP mapping and the current connections for each mapping. You can limit the display to a specific listening port/listening address combination, per routing instance. If you do not specify a routing instance, the default routing instance is assumed.

Traffic between the remote devices and the external systems is expected to be relatively small-sized management requests. Consequently, excessive traffic is not buffered and is dropped by TCP port forwarding. TCP port forwarding does not maintain or recover established TCP connections in the event of a graceful Routing Engine switchover (GRES) or a daemon restart.

You can disable TCP port forwarding by including the **disable** statement at the **[edit system processes]** hierarchy level. You can also configure TCP port forwarding event tracing at the same hierarchy level by including the **traceoptions** statement. See [“Tracing TCP Port Forwarding Events for Troubleshooting” on page 662](#) for more information.

Benefits of TCP Port Forwarding

- Simplifies BNG and remote device configuration and management in topologies that use external management and provisioning systems.

- TCP port forwarding is a generic functionality and can work with any application that can use TCP sessions for communication with remote devices and the BNG.
- Provides several options for tuning the TCP connections to your needs, including restriction to specific IPv4 prefixes, specific listening and forwarding address and port combinations, and the maximum number of allowed connections.

Release History Table

Release	Description
18.3R1	Starting in Junos OS Release 18.3R1, TCP port forwarding (also referred to as TCP forwarding) enables a BNG to mediate communication between its connected access nodes and service provider back-office systems, such as external management and provisioning systems and TACACS+ servers.

RELATED DOCUMENTATION

| [Configure TCP Port Forwarding for Remote Device Management](#) | 659

Configure TCP Port Forwarding for Remote Device Management

To use TCP port forwarding, you configure the mapping between the TCP listening address/listening port combination on the BNG and the TCP port forwarding address/port combination where the BNG forwards the incoming data stream. TCP port forwarding is used when the BNG, together with one or more access nodes, is treated by an external management or provisioning system as a single addressable point of management. The remote devices have private addresses and are not publicly accessible. The TCP port forwarding connections enable the BNG to demultiplex and multiplex management requests exchanged between the access nodes and the management system.

The listening port is monitored by the BNG for connections to be triggered by external management systems or a remote device. The listening address is a particular IPv4 address on the BNG that the triggering entity (external management/provisioning system or remote device) must use when attempting to trigger connections on the listening port.

By default, TCP connections are accepted from any source prefix. You can optionally configure one or more IPv4 prefixes from which TCP connections are accepted on the listening port. You can use a /32 IPv4 mask to specify a single address as the source or you can use other masks to specify an IPv4 subnet as the source. You can configure an unlimited number of prefixes for each listening port. To configure multiple prefixes, however, you must include the statement multiple times, once for each additional source prefix.

NOTE: Although not shown in the following steps, you can also configure TCP port forwarding in a non-default routing instance.

To configure a TCP mapping of a single TCP connection pair for TCP port forwarding:

1. Configure a unique combination of listening port and listening address for each TCP mapping.

```
[edit system services tcp-forwarding]
user@host# set listening-port port-number listening-address ipv4-listening-address
```

2. (Optional) Restrict the IPv4 prefixes from which TCP connections are accepted on the listening port. When you do not configure an allowed source, TCP connections are accepted from any source prefix.

```
[edit system services tcp-forwarding listening-port port-number listening-address ipv4-listening-address]
user@host# set allowed-source ipv4-prefix
```

3. Define the IPv4 address to which MX BNG must open the second connection of the TCP pair after it opens the first connection triggered on the listening port/listening address combination. All packets

received on one connection of the TCP pair are transmitted on the peer (second) connection. This address is used with the forwarding port to open the peer connection.

```
[edit system services tcp-forwarding listening-port port-number listening-address ipv4-listening-address]
user@host# set forwarding-address ipv4-forwarding-address
```

4. Define the TCP port of the peer (second) connection of the TCP pair. This port is used with the forwarding address to open the peer connection.

```
[edit system services tcp-forwarding listening-port port-number listening-address ipv4-listening-address]
user@host# set forwarding-port forwarding-port-number
```

5. (Optional) Set a limit on the number of simultaneous TCP connections that the BNG allows on a single listening port. Connection requests received after this limit is reached are rejected.

```
[edit system services tcp-forwarding listening-port port-number listening-address ipv4-listening-address]
user@host# set max-connections number
```

NOTE: In addition to this per-listening port limit, TCP port forwarding has a system-wide limit of 128 TCP connections (64 connection pairs) across all routing instances and listening ports.

The following sample configuration might be used for the topology shown in [“TCP Port Forwarding for Remote Device Management” on page 655](#). In each step, the listening address is the public address of the BNG for management. A different listening port is assigned for the TACACS+ server, the management platform, and each remote device.

1. Configure the TACACS+ server connection. The BNG monitors port 8020 and its public address for TCP traffic from any of its remote devices to the TACACS server. It accepts traffic only from the subnet shared by the OLTs. It forwards acceptable traffic to the TACACS+ server on the IANA-assigned port number for TACACS, 49. The BNG supports four simultaneous TCP connections on the listening port/address combination, one for each OLT.

```
[edit system services tcp-forwarding]
user@host# edit listening-port 8020 listening-address 203.0.113.50
user@host# set allowed-source 192.0.0.1/24
user@host# set forwarding-address 198.51.100.1
user@host# set forwarding-port 49
user@host# set max-connections 4
```

2. Configure the NETCONF XML protocol connection to each remote device: OLT1, OLT2, OLT3, and OLT4. The BNG monitors its public address and four different ports for TCP traffic from the management platform to the remote devices. Each port is associated with one of the remote devices. The BNG accepts traffic only from the management platform address, 198.51.100.3. Accepted traffic is forwarded to the associated device on the IANA-assigned port number for the NETCONF XML protocol over SSH, 830. Only one TCP connection is supported for each device.

- a. Configure the NETCONF XML protocol connection to OLT1.

```
[edit system services tcp-forwarding]  
user@host# edit listening-port 8000 listening-address 203.0.113.50  
user@host# set allowed-source 198.51.100.3/32  
user@host# set forwarding-address 192.0.0.2  
user@host# set forwarding-port 830  
user@host# set max-connections 1
```

- b. Configure the NETCONF XML protocol connection to OLT2.

```
[edit system services tcp-forwarding]  
user@host# edit listening-port 8001 listening-address 203.0.113.50  
user@host# set allowed-source 198.51.100.3/32  
user@host# set forwarding-address 192.0.0.3  
user@host# set forwarding-port 830  
user@host# set max-connections 1
```

- c. Configure the NETCONF XML protocol connection to OLT3.

```
[edit system services tcp-forwarding]  
user@host# edit listening-port 8002 listening-address 203.0.113.50  
user@host# set allowed-source 198.51.100.3/32  
user@host# set forwarding-address 192.0.0.4  
user@host# set forwarding-port 830  
user@host# set max-connections 1
```

- d. Configure the NETCONF XML protocol connection to OLT4.

```
[edit system services tcp-forwarding]  
user@host# edit listening-port 8003 listening-address 203.0.113.50  
user@host# set allowed-source 198.51.100.3/32  
user@host# set forwarding-address 192.0.0.5  
user@host# set forwarding-port 830  
user@host# set max-connections 1
```


RELATED DOCUMENTATION

TCP Port Forwarding for Remote Device Management | 655

Tracing TCP Port Forwarding Events for Troubleshooting

IN THIS SECTION

- [Configuring the TCP Port Forwarding Trace Log Filename | 663](#)
- [Configuring the Number and Size of TCP Port Forwarding Log Files | 663](#)
- [Configuring Access to the TCP Port Forwarding Log File | 663](#)
- [Configuring a Regular Expression for TCP Port Forwarding Messages to Be Logged | 664](#)
- [Configuring the TCP Port Forwarding Tracing Flags | 664](#)
- [Configuring the Severity Level to Filter Which TCP Port Forwarding Messages Are Logged | 665](#)

The Junos OS trace feature tracks TCP port forwarding operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file located in the `/var/log` directory. By default, the router uses the filename `tcpfwdd`. You can specify a different filename, but you cannot change the directory in which trace files are located.
2. When the trace log file *filename* reaches 128 kilobytes (KB), it is compressed and renamed *filename.0.gz*. Subsequent events are logged in a new file called *filename*, until it reaches capacity again. At this point, *filename.0.gz* is renamed *filename.1.gz* and *filename* is compressed and renamed *filename.0.gz*. This process repeats until the number of archived files reaches the maximum file number. Then the oldest trace file—the one with the highest number—is overwritten.

You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the [System Log Explorer](#).)

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

The following topics describe how to configure all aspects of tracing TCP port forwarding operations:

Configuring the TCP Port Forwarding Trace Log Filename

By default, the name of the file that records trace output for TCP port forwarding is **tcpfwd**. You can specify a different name with the **file** option.

To configure the filename for TCP port forwarding tracing operations:

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

```
[edit system processes tcp-forwarding traceoptions]
user@host# set file tcpfwd_1
```

Configuring the Number and Size of TCP Port Forwarding Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format **.number.gz**. The newest archived file is **.0.gz** and the oldest archived file is **.(maximum number)-1.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, **filename**, reaches 2 MB, **filename** is compressed and renamed **filename.0.gz**, and a new file called **filename** is created. When the new **filename** reaches 2 MB, **filename.0.gz** is renamed **filename.1.gz** and **filename** is compressed and renamed **filename.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, **filename.19.gz**, is simply overwritten when the next oldest file, **filename.18.gz** is compressed and renamed to **filename.19.gz**.

To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output.

```
[edit system processes tcp-forwarding traceoptions]
user@host# set file tcpfwd_1 _logfile_1 files 20 size 2097152
```

Configuring Access to the TCP Port Forwarding Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

To specify that all users can read the log file:

- Configure the log file to be world-readable.

```
[edit system processes tcp-forwarding traceoptions]  
user@host# set file tcpfwd_1 _logfile_1 world-readable
```

To explicitly set the default behavior, only the user who configured tracing can read the log file:

- Configure the log file to be no-world-readable.

```
[edit system processes tcp-forwarding traceoptions]  
user@host# set file tcpfwd_1 _logfile_1 no-world-readable
```

Configuring a Regular Expression for TCP Port Forwarding Messages to Be Logged

By default, the trace operation output includes all messages relevant to the logged events.

You can refine the output by including regular expressions to be matched.

To configure regular expressions to be matched:

- Configure the regular expression.

```
[edit system processes tcp-forwarding traceoptions]  
user@host# set file tcpfwd_1 _logfile_1 match regex
```

Configuring the TCP Port Forwarding Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

```
[edit system processes tcp-forwarding traceoptions]  
user@host# set flag flag-name
```

Configuring the Severity Level to Filter Which TCP Port Forwarding 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 tcp-forwarding traceoptions]  
user@host# set level severity
```

RELATED DOCUMENTATION

| [TCP Port Forwarding for Remote Device Management](#) | 655

Configuring IPFIX Mediation for Remote Device Monitoring

IN THIS CHAPTER

- [IPFIX Mediation on the BNG | 667](#)
- [Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data | 672](#)

IPFIX Mediation on the BNG

Traffic flow is a way of conceptualizing how IP data traffic passes through the various components of your network. A flow consists of a set of IP packets that pass an observation point in the network during a specific time interval. The set is defined by common properties:

- One or more packet, transport, or application header fields
- One or more characteristics of the packet
- One or more fields derived from how the packet is handled

For example, a particular flow might include packets with the same destination IP address and destination port number, number of MPLS labels, next-hop address, and output interface.


The IP Flow Information Export (IPFIX) protocol is a mechanism for transmitting traffic flow information in the form of flow records over your network from an exporting process to a collecting process. Each flow record is generated by a monitoring process and contains information about a specific flow at the observation point, such as the total number of bytes for all packets in the flow and the source IP address. The device that hosts one or more exporting processes is called an exporter or IPFIX device. The device that receives (collects) the flow records from one or more exporting processes is called the collector.

Starting in Junos OS Release 18.3R1, you can configure an MX Series router acting as a BNG to be an intermediary device between IPFIX exporters and collectors. As an IPFIX mediator, the BNG functions as both a collector and an exporter. The IPFIX mediator function collects performance management data via IPFIX records from downstream access network devices such as OLTs and advanced ONUs (with integrated functions such as IPFIX exporter, VOIP SIP client, and so on). This data along with local performance management data from the MX BNG is aggregated and relayed to an upstream IPFIX collector. From the reference point of the IPFIX collector, IPFIX mediation enables the router and its associated access network

devices to appear as a single IPFIX export source leveraging a single TCP/IP connection between the MX BNG and the upstream collector.

Figure 28 on page 668 shows a Passive Optical Network (PON) topology where the BNG IPFIX mediator connected to downstream OLTs, which are in turn connected further downstream to ONTs in residences. The downstream devices export flow information to the mediator over TCP/IP connections; the mediator collects the flow information from the downstream devices. The mediator then processes the flow information and exports it upstream to the IPFIX collector over a TCP or Transport Layer Security (TLS) connection.

Figure 28: Sample Network Topology for IPFIX Mediation



The IPFIX Mediator function enables the BNG and its associated downstream devices to be represented as a single IPFIX exporter towards the IPFIX collector. The data records are not formatted, which optimizes the efficiency of the data stream. A template record, sometimes referred to simply as a template, specifies the semantics and structure for a flow record as an ordered sequence of <type, length> pairs. Template records are sent either before the data records or inline with them.

Each template record includes the template header and one or more field specifiers corresponding to information elements in the data records. The template header includes the template ID and a count of the fields in the template record. The template ID is unique to the transport session and observation domain

(where the traffic flow was observed). Effectively, the ID is unique to the TCP connection between a downstream exporting device and the mediator. Consequently, different downstream devices are likely to use different template IDs for the same record type.

Template ID Reconciliation

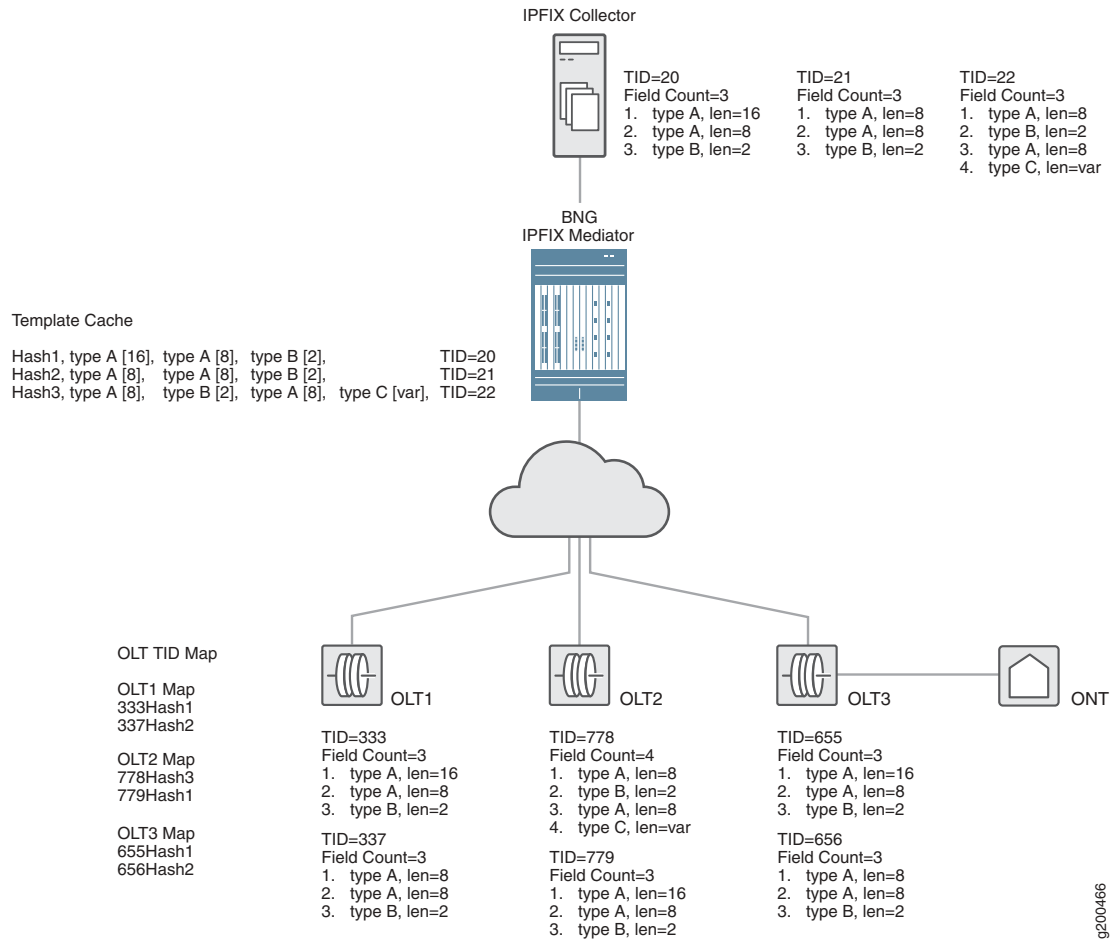
One aspect of the mediation processing is template ID reconciliation. The mediator maintains a cache of unique template records received from the downstream exporters. Matching template records received from different export sources are mapped to the same instance of the record in the template cache. The incoming template records are matched according to the hash value of the number, type, length, and order of the record fields. In other words, the mediator uniquely identifies template records independent of their IDs.

This enables the mediator to assign a new template ID for each unique received template ID. The new upstream ID is used for exporting the template record and data records to the upstream collector. Each new ID is unique to the transport session (TCP or TLS) between the mediator and the collector. This processing results in significantly streamlined communication between the mediator and the collector compared to sending records separately that match except for their template IDs.

[Figure 29 on page 670](#) shows how reconciliation works.

1. The IPFIX mediator receives two template records with different IDs from each OLT.
2. By comparing the hash value for the number and order of the fields and the type and length values for each field, the mediator determines that the six template records from the OLTs represent only three unique records, as follows:
 - The template records with IDs of 333 (OLT1), 779, (OLT2), and 655 (OLT3) all have the same hash value and consequently describe the same record.
 - The template records with IDs of 337 (OLT1) and 656 (OLT3) both have the same hash value and consequently describe the same record.
 - The template record with ID of 778 (OLT2) has a hash value that does not match any other records.
3. Each unique template record is stored in the template cache and assigned a new template ID that is used for sending template and data records to the collector.

Figure 29: Template ID Reconciliation



NOTE: If the IPFIX mediator receives any data records without receiving a corresponding template record in the same TCP session, it discards the data records and logs the event.

The IPFIX mediator functions in a pass-through capacity for the data records from the downstream devices. It does not modify the data records other than changing the template ID for export to the collector. The mediator does not differentiate the data received from different downstream devices; that function is left to the IPFIX collector.

IPFIX Mediation and Network Analytics

IPFIX mediation on the MX Series router employs plug-ins for the **ipfix** network analytics service agent to receive, process, and export IPFIX records. The input plug-in (**input-ipfix**) listens for IPFIX messages on TCP connections from the downstream exporting devices, using port 4739 by default. No other message

types are expected or accepted. The output plug-in (**output-ipfix**) reconciles the received records and sends them to the destination IPFIX collector, which is assumed to listen for them on TCP port 4740 by default. Both plug-ins enable you to configure different parameters for IPFIX mediation. For example, whether the mediator attempts a TLS or TCP connection to the collector is determined by the configuration of certificate options in the output plug-in.

NOTE: The IPFIX plug-ins work only with each other and not with any other analytics plug-in.

Benefits of IPFIX Mediation

- An IPFIX mediator reduces the load on the collector without a loss of information. As the amount of traffic grows in a specific network, the capacity of a single collector to process flow records from multiple exporters can be exceeded. Packet sampling and aggregation can reduce the amount of data to be processed, at the risk of the potential loss of small flows and the detailed information that might be needed to detect and deal with some traffic changes and anomalous behavior.
- An IPFIX mediator provides the flexibility needed when you use multiple traffic monitoring applications. Different applications may require different levels of information, such as packet level versus flow level. These different needs might force the exporter to run different metering tasks to generate flow records, straining limited resources on the device.
- An IPFIX mediator simplifies the accurate monitoring, processing, and exporting of information in networks with a variety of IPFIX devices from multiple vendors, running multiple software releases. A single BNG can mediate the differences between many connected IPFIX devices before exporting flow records to the collector, removing that burden from the individual collectors.

Release History Table

Release	Description
18.3R1	Starting in Junos OS Release 18.3R1, you can configure an MX Series router acting as a BNG to be an intermediary device between IPFIX exporters and collectors.

RELATED DOCUMENTATION

| [Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data](#) | 672

Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data

IPFIX mediation uses the **ipfix** analytics service agent. The service agent uses input and output plug-ins specific to IPFIX. The plug-ins configure aspects of the collecting and exporting functions for the mediator, such as TCP ports and the collector address. The input plug-in takes in the IPFIX flow data from the downstream devices. The output plug-in converts the data to the IPFIX format and exports it to the IPFIX collector. Data conversion is particularly important because users may have a variety of exporting devices using different formats. Converting the formats to a common form on the mediator alleviates the need to have specific collectors for different formats.

Your configuration for the output plug-in determines whether the IPFIX mediator sends records to the collector over a TCP connection or a TLS connection:

- When you configure any of the certificate options (**collector-ca-certificate**, **collector-certificate-key**, or **collector-certificate**), the mediator attempts to make a TLS connection.
- If none of the certificate options is configured, the mediator attempts to make a TCP connection.

To configure IPFIX mediation:

1. Access the IPFIX service agent configuration.

```
[edit services analytics agent]
user@host# edit service-agents ipfix
```

2. Configure parameters for the IPFIX input plug-in.

```
[edit services analytics agent service-agents ipfix]
user@host# edit inputs input-ipfix
```

NOTE: Although each of the parameters has a default value, you must configure at least one of the parameters to enable the plug-in. If you configure only one parameter and want to use the default value, you must specify that value.

- a. (Optional) Specify the maximum number of TCP connections that the IPFIX mediator can have. The default value is 100.

```
[edit services analytics agent service-agents ipfix inputs input-ipfix]
user@host# set parameters maximum-connections number
```

- b. (Optional) Specify the TCP port that the IPFIX mediator uses to receive TCP packets from the downstream devices. The default value is 4739.

```
[edit services analytics agent service-agents ipfix inputs input-ipfix]
user@host# set parameters tcp-port port-number
```

- c. Specify the name of the VRF (routing instance) where IPFIX packets are accepted from the downstream devices.

```
[edit services analytics agent service-agents ipfix inputs input-ipfix]
user@host# set parameters vrf-name name
```

3. Configure parameters for the output plug-in.

```
[edit services analytics agent service-agents ipfix]
user@host# edit outputs output-ipfix
```

- a. Specify the IP address of the upstream IPFIX collector. This is a mandatory option.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-address ip-address
```

- b. (Optional) Specify the path for the certificate that is used to sign the peer certificate at the peer (IPFIX collector) level. The certificate is provided by a trusted certificate authority (CA) and is expected to be in .pem container format.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-ca-certificate file-path
```

- c. (Optional) Specify the path for the client certificate that the server (IPFIX collector) uses to authenticate the client and to enable mutual authentication. The fully-qualified domain name (FQDN) of both the client and the server are stored in the certificate's Subject Alternative Name field when the client and server certificates are generated. The certificate is expected to be in .pem container format.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-certificate file-path
```

- d. (Optional) Specify the path of the private key file that is loaded to decrypt the encrypted message sent from the peer.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-certificate-key file-path
```

- e. (Optional) Specify how many seconds the output plug-in waits before retrying the connection to the IPFIX collector. The default value is 20.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-connection-retry-interval seconds
```

- f. (Optional) Specify the TCP port that the IPFIX mediator uses to connect to the IPFIX collector. The default value is 4740.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-tcp-port port-number
```

- g. (Optional) Specify the name of the VRF (routing instance) in which IPFIX packets are routed to the IPFIX collector. The default value is **default**.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-vrf-name vrf-name
```

In the following sample configuration, the input plug-in is configured so that the IPFIX mediator accepts up to 125 TCP connections from its downstream devices. Records are accepted in the RI-ipfix-1 routing instance. The TCP port is not configured, so the plug-in listens on the default port, 4739.

```
[edit services analytics agent service-agents ipfix]
user@host# set inputs input-ipfix parameters maximum-connections 125
user@host# set inputs input-ipfix parameters vrf-name RI-ipfix-1
```

The following example configuration for the output plug-in specifies that:

- Records are exported to the collector at 198.51.100.200.
- If the connection to the collector is not successful, the plug-in attempts to make the connection at 15-second intervals.
- The configuration includes paths for collector certificates, so the export connection is over TLS rather than TCP.
- The TCP port is not configured, so the collector is expected to listen on the default port, 4740.
- No routing instance is configured for the collector, so it accepts packets in the default routing instance.

```
user@host# edit services analytics agent service-agents ipfix
user@host# set outputs output-ipfix parameters collector-address 198.51.100.200
user@host# set outputs output-ipfix parameters collector-ca-certificate /var/tmp/ca.pem
user@host# set outputs output-ipfix parameters collector-certificate /var/tmp/client.pem
```

```
user@host# set outputs output-ipfix parameters collector-certificate-key /var/tmp/example.com.key  
user@host# set outputs output-ipfix parameters collector-connection-retry-interval 15
```

RELATED DOCUMENTATION

| [IPFIX Mediation on the BNG](#) | 667

Collection and Export of Local Telemetry Data on the IPFIX Mediator

IN THIS CHAPTER

- Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector | 677
- Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator | 680

Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector

Starting in Junos OS Release 18.4R1, the input-jti-ipfix plug-in for the IPFIX service agent collects telemetry data from the local Junos Telemetry Interface (JTI) on a BNG configured as an IPFIX mediator. The output-ipfix plug-in translates the gRPC data received from a nonconfigurable set of well-defined sensor types into specific IPFIX records. You configure a record group for the input plug-in that consists of one or more of the predefined IPFIX records. Each record is associated with a specific set of telemetry sensors on the BNG, as listed in [Table 53 on page 677](#).

Table 53: IPFIX Records and Associated Telemetry Sensors (gRPC Path)

IPFIX Record Name	Sensors Collected by the Record
address-pool-utilization	/junos/system/subscriber-management/aaa/address-assignment-statistics/logical-system-routing-instanc
chassis-inventory	<div>/components/component[name='Routing EngineX']/properties/property[name='fru-model-number'] /components/component[name='FPCx']/properties/property[name='fru-model-number'] /components/component[name='Routing Engine0']/state/id /components/component[name='FPCx']/state/id /components/component[name='Routing EngineX']/properties/property[name='hardware-rev']/ /components/component[name='FPCx']/properties/property[name='hardware-rev']/ /components/component[name='FPCx']/properties/property[name='software-rev']/</div>

Table 53: IPFIX Records and Associated Telemetry Sensors (gRPC Path) (continued)

IPFIX Record Name	Sensors Collected by the Record
chassis-power	/components/component[name='Chassis']/properties/property[name='power-system-capacity']/ /components/component[name='Chassis']/properties/property[name='power-system-remaining']/
dhcpv4-server-statistics	/junos/system/subscriber-management/client-protocols/dhcp/v4/routing-instances/routing-instance/ser
interface-meta-data	/junos/system/subscriber-management/dynamic-interfaces/interfaces/meta-data/interface
port-statistics	/interfaces/interface/state/counters/
resource-utilization	/components/component[name='Routing Engine0']/properties/property[name='memory-dram-used']/ /components/component[name='Routing Engine1']/properties/property[name='memory-dram-used']/ /components/component[name='Routing Engine0']/properties/property[name='memory-dram-installed']/ /components/component[name='Routing Engine1']/properties/property[name='memory-dram-installed']/ /components/component[name='FPCx']/properties/property[name='memory-utilization-heap']/ /compon Engine0']/properties/property[name='memory-utilization-buffer']/ /components/component[name='Routing Engine1']/properties/property[name='memory-utilization-buffe /components/component[name='FPCx']/properties/property[name='memory-utilization-buffer']/ /compo Engine0']/properties/property[name='cpu-utilization-idle']/ /components/component[name='Routing Engine1']/properties/property[name='cpu-utilization-idle']/
subscriber-statistics	/junos/system/subscriber-management/dynamic-interfaces/interfaces/subscriber-statistics/interface
thermal	/components/component[name='Chassis']/properties/property[name='temperature-ambient'] /components/component[name='RoutingEngine0']/properties/property[name='temperature']/state/value /components/component[name='RoutingEngine1']/properties/property[name='temperature']/state/value /components/component[name='FPCx']/properties/property[name='temperature-exhaust-x']/state/value
uptime	/components/component[name='FPCx']/properties/property[name='uptime']/

You can configure the frequency with which data is collected and reported to an IPFIX collector. The reporting interval has a default value, but some telemetry data, such as subscriber statistics, is more dynamic than other data, such as chassis temperature. A shorter reporting interval may be more useful for the more dynamic data. You can configure the reporting interval for the record group, but not for individual records.

The template IDs for the translated gRPC data are drawn from the same template ID space as the IPFIX mediator. Template IDs in the range 256 though 400 are reserved for the translation of telemetry data.

NOTE: For more information about sensors, see *Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)*.

For detailed information about the Juniper Telemetry Interface, see *Junos Telemetry Interface User Guide*.

Benefits of Telemetry Data Collection

- Leverages the IPFIX mediation structure to collect data about hardware, resources, and user statistics for better remote management of the BNG and subscribers.

Release History Table

Release	Description
18.4R1	Starting in Junos OS Release 18.4R1, the input-jti-ipfix plug-in for the IPFIX service agent collects telemetry data from the local Junos Telemetry Interface (JTI) on a BNG configured as an IPFIX mediator.

RELATED DOCUMENTATION

Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator 680
Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data 672
IPFIX Mediation on the BNG 667
<i>Understanding Enabling Export of Subscriber Statistics and Queue Statistics for Dynamic Interfaces and Interface-Sets</i>

Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator

You can configure the input-jti-ipfix plug-in for the IPFIX service agent to collect telemetry (gRPC) data from the local Junos Telemetry Interface (JTI) on a BNG configured as an IPFIX mediator. In addition to streaming IPFIX records from the input-ipfix plug-in, the output-ipfix plug-in also translates the gRPC data received from the input-jti-ipfix plug-in into corresponding IPFIX data records.

You configure a record group for the input-jti-ipfix plug-in that consists of one or more predefined IPFIX records. Each predefined record is associated with a specific, nonconfigurable set of telemetry sensors on the BNG. You can configure the frequency at which the sensor records are exported to an IPFIX collector; the IPFIX collector is configured with the output-ipfix plug-in.

Before you begin, you must enable the IPFIX service agent by configuring at least one parameter for the input-ipfix plug-in.

To configure local telemetry data collection and reporting:

1. Access the IPFIX service agent configuration.

```
[edit services analytics agent]
user@host# edit service-agents ipfix
```

2. Configure parameters for the IPFIX telemetry input plug-in.

```
[edit services analytics agent service-agents ipfix]
user@host# edit inputs input-jti-ipfix
```

- a. Specify the name of a group of records to collect telemetry data.

```
[edit services analytics agent service-agents ipfix inputs input-jti-ipfix]
user@host# edit parameters record-group group-name
```

- b. Specify the record that you want to add to the group.

```
[edit services analytics agent service-agents ipfix inputs input-jti-ipfix parameters record-group group-name]
user@host# set record ipfix-record-name
```

- c. (Optional) Configure a reporting interval for the record group when you do not want to use the default value (900 seconds).

```
[edit services analytics agent service-agents ipfix inputs input-jti-ipfix parameters record-group group-name]
```

```
user@host# set reporting-interval seconds
```

3. Configure parameters for the IPFIX output plug-in. This is the same configuration you use when you configure the IPFIX mediation.

```
[edit services analytics agent service-agents ipfix]
user@host# edit outputs output-ipfix
```

- a. Specify the IP address of the upstream IPFIX collector. This is a mandatory option.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-address ip-address
```

- b. (Optional) Specify the path for the certificate that is used to sign the peer certificate at the peer (IPFIX collector) level. The certificate is provided by a trusted certificate authority (CA) and is expected to be in .pem container format.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-ca-certificate file-path
```

- c. (Optional) Specify the path for the client certificate that the server (IPFIX collector) uses to authenticate the client and to enable mutual authentication. The fully-qualified domain name (FQDN) of both the client and the server are stored in the certificate's Subject Alternative Name field when the client and server certificates are generated. The certificate is expected to be in .pem container format.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-certificate file-path
```

- d. (Optional) Specify the path of the private key file that is loaded to decrypt the encrypted message sent from the peer.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-certificate-key file-path
```

- e. (Optional) Specify how many seconds the output plug-in waits before retrying the connection to the IPFIX collector. The default value is 20.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-connection-retry-interval seconds
```

- f. (Optional) Specify the TCP port that the IPFIX mediator uses to connect to the IPFIX collector. The default value is 4740.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-tcp-port port-number
```

- g. (Optional) Specify the name of the VRF (routing instance) in which IPFIX packets are routed to the IPFIX collector. The default value is **default**.

```
[edit services analytics agent service-agents ipfix outputs output-ipfix]
user@host# set parameters collector-vrf-name vrf-name
```

The following sample configuration includes three record groups for the telemetry input plug-in, high-frequency, baseline, and background:

- The high-frequency group subscribes to the subscriber-statistics and port-statistics record. Because statistics data is dynamic and changes frequently, the reporting interval is set to five minutes, which is much less than the default.
- The baseline record group subscribes to the address-pool-utilization and dhcpv4-server-statistics records; the reporting interval is left at the default value, 15 minutes.
- The background record group subscribes to the thermal and chassis-inventory records. These probably do not change frequently, so the reporting interval is set to six hours.

To enable the IPFIX plug-in, you must configure at least one parameter; in this example, the maximum number of TCP connections is set to 200.

Finally, the IP address and listening port for the IPFIX collector is configured in the output plug-in.

```
[edit services analytics agent service-agents ipfix]
inputs input-jti-ipfix {
  parameters {
    record-group high-frequency {
      record subscriber-statistics;
      record port-statistics;
      reporting-interval 300;
    }
    record-group baseline {
      record address-pool-utilization;
      record dhcpv4-server-statistics;
    }
    record-group background {
      record thermal;
      record chassis-inventory;
```

```
        reporting-interval 21600;
    }
}
inputs input-ipfix {
    parameters {
        maximum-connections 200;
    }
}
outputs output-ipfix {
    parameters {
        collector-address 192.0.2.2;
        collector-port 6589;
    }
}
}
```

You can use the [show services analytics agent](#) command to display information about the service agents.

RELATED DOCUMENTATION

[Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector | 677](#)

[Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data | 672](#)

[IPFIX Mediation on the BNG | 667](#)

9

PART

Troubleshooting

[Contacting Juniper Networks Technical Support](#) | **687**

Contacting Juniper Networks Technical Support

IN THIS CHAPTER

- [Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support | 687](#)

Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support

Problem

Description: When you experience a subscriber access problem in your network, we recommend that you collect certain logs before you contact Juniper Networks Technical Support. This topic shows you the most useful logs for a variety of network implementations. In addition to the relevant log information, you must also collect standard troubleshooting information and send it to Juniper Networks Technical Support in your request for assistance.

Solution

To collect standard troubleshooting information:

- Redirect the command output to a file.

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

To configure logging to assist Juniper Networks Technical Support:

1. Review the following blocks of statements to determine which apply to your configuration.

[edit]

```

set system syslog archive size 100m files 25
set system auto-configuration traceoptions file filename
set system auto-configuration traceoptions file filename size 100m files 25
set protocols ppp-service traceoptions file filename size 100m files 25
set protocols ppp-service traceoptions level all
set protocols ppp-service traceoptions flag all
set protocols ppp traceoptions file filename size 100m files 25
set protocols ppp traceoptions level all
set protocols ppp traceoptions flag all
set protocols ppp monitor-session all
set interfaces pp0 traceoptions flag all
set demux traceoptions file filename size 100m files 25
set demux traceoptions level all
set demux traceoptions flag all
set system processes dhcp-service traceoptions file filename
set system processes dhcp-service traceoptions file size 100m
set system processes dhcp-service traceoptions file files 25
set system processes dhcp-service traceoptions flag all
set class-of-service traceoptions file filename
set class-of-service traceoptions file size 100m
set class-of-service traceoptions flag all
set class-of-service traceoptions file files 25
set routing-options traceoptions file filename
set routing-options traceoptions file size 100m
set routing-options traceoptions flag all
set routing-options traceoptions file files 25
set interfaces traceoptions file filename
set interfaces traceoptions file size 100m
set interfaces traceoptions flag all
set interfaces traceoptions file files 25
set system processes general-authentication-service traceoptions file filename
set system processes general-authentication-service traceoptions file size 100m
set system processes general-authentication-service traceoptions flag all
set system processes general-authentication-service traceoptions file files 25

```


2. Copy the relevant statements into a text file and modify the log filenames as you want.
3. Copy the statements from the text file and paste them into the CLI on your router to configure logging.
4. Commit the logging configuration to begin collecting information.

NOTE: The maximum file size for DHCP local server and DHCP relay log files is 1 GB. The maximum number of log files for DHCP local server and DHCP relay is 1000.

BEST PRACTICE: Enable these logs only to collect information when troubleshooting specific problems. Enabling these logs during normal operations can result in reduced system performance.

RELATED DOCUMENTATION

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

10

PART

Configuration Statements and Operational Commands

Configuration Statements | **693**

Operational Commands | **1253**

Configuration Statements

IN THIS CHAPTER

- access-domain (Remote Device Management) | 703
- accounting (Dynamic IGMP Interface) | 705
- accounting (Dynamic MLD Interface) | 706
- action | 707
- address (LRF Profile) | 708
- address-mapping (Application Identification) | 709
- adf (Dynamic Firewalls) | 710
- adjustment-control-profiles | 712
- adjust-minimum (Dynamic Shaping and Scheduling) | 713
- adjust-percent (Dynamic Schedulers) | 714
- agent (Analytics) | 715
- aggregate (Hierarchical Policier) | 718
- alt-name (Application Identification) | 719
- analytics | 720
- ancp (Adjustment Control Profiles) | 727
- application (Adjustment Control Profiles) | 729
- application (Application Identification) | 730
- application-groups (PCC Rules) | 732
- application-identification (Application Identification) | 734
- application-identification-profile (Service Set) | 736
- applications (PCC Rules) | 737
- apply-groups (Subscriber Secure Policy) | 739
- apply-groups-except (Subscriber Secure Policy) | 740
- authentication-order | 741
- bandwidth (Tunnel Services) | 743
- bandwidth-limit (Policier) | 745
- bandwidth-percent | 747
- buffer-size (Dynamic Scheduling) | 750

- burst-size-limit (Hierarchical Policer) | 752
- burst-size-limit (Policer) | 754
- bytes (Dynamic Traffic Shaping) | 756
- cacheable (Application Identification) | 757
- captive-portal-content-delivery (Captive Portal Content Delivery) | 758
- captive-portal-content-delivery-profile (Services) | 761
- cell-mode (Dynamic Traffic Shaping) | 763
- chain-order (Application Identification) | 765
- check-bytes (Application Identification) | 766
- class (Defining Login Classes) | 767
- class-of-service (Dynamic Profiles) | 778
- classifiers (Dynamic CoS Application) | 780
- code (Application Identification) | 781
- collector (LRF Profile) | 782
- collector (LRF Rule) | 783
- color-aware | 784
- color-blind | 786
- committed-burst-size | 788
- committed-information-rate | 790
- compatibility (Application Identification) | 792
- connection-limit | 793
- context (Application Identification) | 795
- delay-buffer-rate (Dynamic Traffic Shaping) | 797
- description (Application Identification) | 798
- destination (Application Identification) | 799
- destination (LRF Profile) | 800
- destination-address (Subscriber Secure Policy) | 801
- destination-port (Subscriber Secure Policy) | 802
- ddos-protection (DDoS) | 803
- dhcp-tags (Adjustment Control Profiles) | 807
- direction (Application Identification) | 809
- direction (Service Data Flow Filters) | 810
- disable (Dynamic IGMP) | 811
- disable (Dynamic MLD) | 812

- download (Application Identification) | 813
- drop-policy (Subscriber Secure Policy) | 814
- drop-profile (Dynamic Schedulers) | 815
- drop-profile-map (Dynamic Schedulers) | 817
- dscp (Dynamic Classifiers) | 818
- dscp (Dynamic Rewrite Rules) | 819
- dscp (Subscriber Secure Policy) | 820
- dscp-ipv6 (Dynamic Classifiers) | 821
- dscp-ipv6 (Dynamic Rewrite Rules) | 822
- dtcp-only (System Services) | 823
- dynamic-class-of-service-options (Dynamic Traffic Shaping) | 824
- dynamic-profiles | 826
- effective-shaping-rate | 840
- enable-performance-mode (Application Identification) | 841
- enhanced-mode | 842
- enhanced-mode-override | 845
- enhanced-policer | 846
- excess-burst-size | 847
- excess-priority (Dynamic Schedulers) | 849
- excess-rate (Dynamic Schedulers) | 850
- excess-rate (Dynamic Traffic Shaping) | 852
- excess-rate-high (Dynamic Traffic Shaping) | 854
- excess-rate-low (Dynamic Traffic Shaping) | 856
- exclude (Dynamic MLD Interface) | 857
- fail-filter (Dynamic Profiles) | 858
- family (Dynamic Firewalls) | 859
- family (Dynamic Standard Interface) | 861
- fast-update-filter (Dynamic Firewalls) | 864
- filter (Configuring) | 866
- filter (Dynamic Profiles Filter Attachment) | 868
- filter (Dynamic Profiles Filter Creation) | 870
- filter (Dynamic Interface Unit) | 872
- filter-specific | 874
- firewall (Dynamic Firewalls) | 876

- flow-descriptions | **879**
- flow-tap | **881**
- flow-tap-dtcp | **883**
- flows (PCC Rules) | **885**
- format (LRF Profile) | **886**
- forwarding-class (Dynamic Scheduler Maps) | **887**
- forwarding-class (PCC Action Profiles) | **888**
- forwarding-class (Subscriber Secure Policy) | **889**
- fpc (MX Series 5G Universal Routing Platforms) | **890**
- frame-mode (Dynamic Traffic Shaping) | **892**
- from (Captive Portal Content Delivery Tags) | **894**
- from (PCC Rules) | **895**
- from (Subscriber Secure Policy) | **897**
- gate-status | **898**
- group (Dynamic IGMP Interface) | **900**
- group (Dynamic MLD Interface) | **902**
- group-count (Dynamic MLD Interface) | **903**
- group-increment (Dynamic MLD Interface) | **904**
- group-limit (Dynamic IGMP Interface) | **905**
- group-limit (Dynamic MLD Interface) | **906**
- group-policy (Dynamic IGMP Interface) | **907**
- group-policy (Dynamic MLD Interface) | **908**
- guaranteed-rate (Dynamic Traffic Shaping) | **909**
- hierarchical-policer | **911**
- hierarchical-scheduler (Subscriber Interfaces on MX Series Routers) | **914**
- http-log-multiple-transactions (LRF Profile) | **916**
- icmp-mapping (Application Identification) | **917**
- ieee-802.1 (Dynamic Classifiers) | **918**
- ieee-802.1 (Dynamic Rewrite Rules) | **919**
- if-exceeding (Hierarchical Policer) | **920**
- if-exceeding (Policer) | **922**
- igmp (Dynamic Profiles) | **924**
- immediate-leave (Dynamic IGMP Interface) | **926**
- immediate-leave (Dynamic MLD Interface) | **927**

- inet (Subscriber Secure Policy) | 929
- inet-precedence (Dynamic Classifiers) | 930
- inet-precedence (Dynamic Rewrite Rules) | 931
- inet6 (Subscriber Secure Policy) | 932
- input (Dynamic Service Sets) | 933
- inputs (Analytics) | 934
- interface (Dynamic IGMP) | 938
- interface (Dynamic Interface Sets) | 940
- interface (Dynamic MLD) | 942
- interface (Dynamic Routing Options) | 944
- interface-service (Services Interfaces) | 945
- interface-set (Dynamic Profiles) | 946
- interface-shared | 948
- interface-specific (Dynamic Firewalls) | 949
- interfaces (Dynamic CoS Definition) | 950
- interfaces (Static and Dynamic Subscribers) | 952
- ip-protocol-mapping (Application Identification) | 959
- ipv4-address (Steering Path) | 960
- ipv6-address (Steering Path) | 961
- keep-existing-steering | 962
- local-port-range | 963
- local-ports | 965
- logging-rule (PCC Action Profile) | 967
- logical-bandwidth-policer | 968
- logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces) | 969
- logical-interface-policer | 970
- logical-system (Subscriber Secure Policy) | 973
- login | 974
- loss-priority (Dynamic Schedulers) | 979
- loss-priority high then discard (Three-Color Policer) | 980
- max-queues-per-interface | 982
- match-order (Dynamic Firewalls) | 984
- maximum-bit-rate (PCC Action Profiles) | 985
- member (Application Identification) | 987

- mld (Dynamic Profiles) | **988**
- multicast (Dynamic Routing Options) | **990**
- multicast-interception (Subscriber Secure Policy) | **991**
- netconf (Remote Device Management) | **992**
- no-accounting | **995**
- no-qos-adjust (Dynamic Routing Options) | **996**
- oif-map (Dynamic IGMP Interface) | **997**
- oif-map (Dynamic MLD Interface) | **998**
- order (Application Identification) | **999**
- order-priority (Application Identification) | **1000**
- output (Dynamic Service Sets) | **1001**
- outputs (Analytics) | **1002**
- output-traffic-control-profile (Dynamic CoS Definition) | **1006**
- overhead-accounting (Dynamic Traffic Shaping) | **1007**
- passive (Dynamic IGMP Interface) | **1008**
- passive (Dynamic MLD Interface) | **1009**
- path (Steering) | **1010**
- pattern (Application Identification) | **1011**
- pcc-action-profile (PCC Rules) | **1012**
- pcc-action-profiles | **1014**
- pcc-context | **1016**
- pcc-rule | **1018**
- pcc-rulebases (PCEF) | **1020**
- pcc-rulebases (PCEF Profile) | **1022**
- pcc-rules (PCEF) | **1024**
- pcc-rules (PCEF Profile) | **1026**
- pcef (Dynamic Profiles) | **1028**
- pcef-profile (Service Set) | **1029**
- peak-burst-size | **1030**
- peak-information-rate | **1032**
- physical-interface-policer | **1034**
- policer (Configuring) | **1036**
- policy (Subscriber Secure Policy) | **1038**
- policy-based-logging (LRF Profile) | **1040**

- policy-options (Dynamic Profiles) | **1041**
- policy-statement | **1042**
- port (LRF Profile) | **1048**
- port-range (Application Identification) | **1049**
- post-service-filter (Dynamic Service Sets) | **1050**
- pppoe-tags (Adjustment Control Profiles) | **1051**
- precedence | **1053**
- premium (Hierarchical Policer) | **1055**
- priority (Dynamic Schedulers) | **1057**
- priority (Application Identification With Next Gen Services) | **1058**
- profile (Access) | **1059**
- profile (Captive Portal Content Delivery) | **1066**
- profile (LRF) | **1068**
- profile (Services PCEF) | **1070**
- profiles (PCEF) | **1071**
- profile-type (Dynamic Service Profiles) | **1073**
- promiscuous-mode (Dynamic IGMP Interface) | **1074**
- protocol (Application Identification) | **1075**
- protocol (Dynamic Schedulers) | **1076**
- protocol (Flow Descriptions) | **1077**
- protocol (Subscriber Secure Policy) | **1078**
- protocols (DDoS) | **1079**
- protocols (Dynamic Profiles) | **1091**
- provisioning-method (Remote Device Management) | **1094**
- radius (Access Profile) | **1095**
- radius-coa (Adjustment Control Profiles) | **1099**
- radius-flow-tap | **1101**
- radius-server | **1103**
- rate-limit | **1109**
- rebalance-periodic (Aggregated Ethernet Subscriber Interfaces) | **1110**
- redirect (PCC Action Profiles) | **1111**
- remote-address | **1113**
- remote-device-management | **1115**
- remote-port-range | **1117**

- remote-ports | **1119**
- report (LRF Rule) | **1121**
- rewrite-rules (Dynamic CoS Interfaces) | **1122**
- routing-engine-services | **1123**
- routing-options (Dynamic Profiles) | **1124**
- routing-instance (Subscriber Secure Policy) | **1126**
- routing-instance (PCC Action Profiles) | **1127**
- rpf-check (Dynamic Profiles) | **1129**
- rule (Captive Portal Content Delivery) | **1130**
- rule (LRF) | **1132**
- rule-set (Captive Portal Content Delivery) | **1133**
- scheduler (Dynamic Scheduler Maps) | **1134**
- scheduler-map (Dynamic Traffic Shaping) | **1135**
- scheduler-maps (Dynamic CoS Definition) | **1136**
- schedulers (Dynamic CoS Definition) | **1137**
- service (Dynamic Profiles) | **1138**
- service (Dynamic Service Sets) | **1139**
- service-agents (Analytics) | **1141**
- service-device (Remote Device Management) | **1143**
- service-filter (Dynamic Service Sets) | **1146**
- service-interface (Services Interfaces) | **1147**
- service-set (Application-Aware Control Policy) | **1148**
- service-set (Dynamic Service Sets) | **1150**
- services (Captive Portal Content Delivery) | **1152**
- session-options | **1154**
- shaping-rate (Dynamic Traffic Shaping and Scheduling) | **1158**
- shared-name | **1160**
- signature (Application Identification) | **1161**
- single-rate | **1162**
- snmp (Subscriber Secure Policy) | **1163**
- source (Application Identification) | **1164**
- source (Dynamic IGMP Interface) | **1165**
- source (Dynamic MLD Interface) | **1166**
- source-address (Subscriber Secure Policy) | **1167**

- source-address (LRF Profile) | **1168**
- source-count (Dynamic MLD Interface) | **1169**
- source-increment (Dynamic MLD Interface) | **1170**
- source-ipv4-address | **1171**
- source-port (Subscriber Secure Policy) | **1172**
- ssh (System Services) | **1173**
- ssm-map (Dynamic IGMP Interface) | **1175**
- ssm-map (Dynamic MLD Interface) | **1177**
- ssm-map-policy (Dynamic IGMP Interface) | **1179**
- ssm-map-policy (Dynamic MLD Interface) | **1181**
- stacked-interface-set (Dynamic Profiles) | **1183**
- static (Dynamic IGMP Interface) | **1185**
- static (Dynamic MLD Interface) | **1186**
- static-policy-control | **1187**
- steering | **1189**
- subscriber-leave-timer | **1191**
- tags (Application Identification) | **1192**
- targeted-distribution (Dynamic Demux Interfaces over Aggregated Ethernet) | **1193**
- targeted-distribution (Static Interfaces over Aggregated Ethernet) | **1194**
- tcp-forwarding (Processes) | **1195**
- tcp-forwarding (Remote Device Management) | **1196**
- template (LRF Profile) | **1199**
- template (LRF Rule) | **1200**
- template-tx-interval (LRF Profile) | **1201**
- template-type (LRF Profile) | **1202**
- term (Captive Portal Content Delivery) | **1204**
- term (Dynamic Profiles) | **1206**
- then (Captive Portal Content Delivery) | **1208**
- then (LRF rule) | **1210**
- then (PCC Rules) | **1211**
- three-color-policer (Configuring) | **1213**
- time-limit (LRF Rule) | **1215**
- traceoptions (Analytics Agent) | **1216**
- traceoptions (Captive Portal Content Delivery) | **1218**

- [traceoptions \(TCP Port Forwarding\) | 1220](#)
- [traffic-control-profiles \(Dynamic CoS Definition\) | 1222](#)
- [transmit-rate \(Dynamic Schedulers\) | 1224](#)
- [trigger-type \(LRF Profile\) | 1226](#)
- [tunnel-services \(Chassis\) | 1227](#)
- [two-rate | 1229](#)
- [type \(Application Identification\) | 1230](#)
- [type \(ICMP Mapping for Application Identification\) | 1231](#)
- [uid \(Dynamic Profiles\) | 1232](#)
- [uid-reference | 1233](#)
- [unit \(Dynamic Profiles Standard Interface\) | 1234](#)
- [unit \(Dynamic Traffic Shaping\) | 1239](#)
- [url | 1241](#)
- [user \(Access\) | 1243](#)
- [vendor-support | 1246](#)
- [version \(Dynamic IGMP Interface\) | 1247](#)
- [version \(Dynamic MLD Interface\) | 1248](#)
- [vlan-tag \(Dynamic Classifiers\) | 1249](#)
- [vlan-tag \(Dynamic Rewrite Rules\) | 1250](#)
- [volume-limit \(LRF Rule\) | 1251](#)

access-domain (Remote Device Management)

Syntax

```
access-domain {  
    vlan-id-list [vlan-id-low-vlan-id-high vlan-id]  
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name system services remote-device-management service-device  
    device-name],  
[edit system services remote-device-management service-device device-name]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Specify the Layer 2 domain that is served by the service device. The access domain represents the set of subscriber-facing Layer 2 locations that map to the device. You define the domain when you map subscribers to the device by configuring one or more non-overlapping outer VLAN ID ranges or discrete VLAN IDs that are served by the remote device. All BNG subscriber sessions with outer VLAN tags matching this list are connected to this remote device in the access network. During subscriber negotiation and provisioning, the BNG provisions remote services sourced by the policy server (PCRF) on this remote device.

You can configure the VLAN IDs as discrete values, even if they are consecutive; however, it is more practical to configure ranges of IDs with additional discrete VLAN IDs as needed. You can configure any number of ranges or discrete IDs; no limit is enforced.

The list of VLAN IDs can overlap across remote devices to support redundancy in the access network.

NOTE: To specify more than one range or address in the list, place the values inside square brackets, separated by spaces. For example:

```
[100 102 110 120-130]
```

NOTE: You do not have to specify an access domain (list of VLAN IDs) when you configure a service device. This means that a service device without the access domain can be connected and dynamically assigned an IP address. Then at some later point when the device needs to support subscriber connections, you can configure one or more VLAN IDs or ranges with the **vlan-id-list** statement.

NOTE: You can add new VLAN IDs or ranges to the access domain at any time, even when the device has active subscriber services mapped to it. You cannot delete IDs or ranges when at least one active subscriber service is configured for that discrete VLAN ID or any VLAN ID in the range being deleted.

You can use the [request services remote-device-management reconfigure service-device](#) command to reconfigure a remote device to provision all active subscriber services matching the access domain.

Options

vlan-id-low—(Optional) Lowest VLAN ID in the range,

vlan-id-high—(Optional) Highest VLAN ID in the range.

vlan-id—(Optional) Discrete VLAN ID that is not part of a range.

Required Privilege Level

system—To view this statement in the configuration.

RELATED DOCUMENTATION

[Configuring Remote Device Management for Service Provisioning](#) | 648

[Remote Device Services Manager \(RDSM\) Overview](#) | 631

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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Recording IGMP Join and Leave Events

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

[Dynamic MLD Configuration Overview | 401](#)

[Dynamic IGMP Configuration Overview | 393](#)

Example: Recording MLD Join and Leave Events

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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](#) | [980](#)

address (LRF Profile)

Syntax

```
address collector-address;
```

Hierarchy Level

```
[edit services lrf profile profile-name collector collector-name destination]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Specify the destination IP address of the collector.

Options

collector-address—IP address of the collector.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers](#) | [460](#)

Configuring Logging and Reporting for Junos OS Subscriber Aware

address-mapping (Application Identification)

Syntax

```
address-mapping name {
  destination {
    ip ip-address-prefix;
  }
  source {
    ip ip-address-prefix;
  }
  order order;
  order-priority (high | low);
}
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Define an application signature based on the source or destination IP address.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Options

name—Name given to the application associated with the source or destination IP address.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

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

[Understanding Dynamic Firewall Filters | 226](#)

[Classic Filters Overview | 229](#)

[Basic Classic Filter Syntax | 232](#)

[Guidelines for Configuring Service Filters](#)

adjustment-control-profiles

Syntax

```
adjustment-control-profiles {
  profile-name {
    application {
      (anyp | dhcp-tags | pppoe-tags | radius-coa)
      priority priority;
      algorithm algorithm;
    }
  }
}
```

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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](#) | 183

[Configuring CoS Adjustment Control Profiles](#) | 186

[Verifying the CoS Adjustment Control Profile Configuration](#) | 188

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](#)

[Configuring a Dynamic Shaping-Rate Adjustment for Queues](#)

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*

agent (Analytics)

Syntax

```
agent {
  service-agents {
    agent-name {
      inputs {
        analytics {
          parameters {
            generate-tags value;
            sample-frequency value;
            sensors file-path;
          }
        }
      }
      input-ipfix {
        parameters {
          maximum-connections number;
          tcp-port port-number;
          vrf-name name;
        }
      }
      input-jti-ipfix {
        parameters {
          record-group group-name {
            record ipfix-record-name;
            reporting-interval seconds;
          }
        }
      }
    }
  }
  outputs {
    file {
      parameters {
        path file-path;
      }
    }
    kafka {
      parameters {
        server ip-address;
        topic topic-name;
        encoding encoding-type;
      }
    }
    output-ipfix {
```


RELATED DOCUMENTATION

[Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data | 672](#)

[Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator | 680](#)

Configuring NTF Agent

[IPFIX Mediation on the BNG | 667](#)

[Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector | 677](#)

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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 \(Hierarchical Policer\)](#)

[burst-size-limit \(Hierarchical Policer\) | 752](#)

[hierarchical-policer | 911](#)

[if-exceeding \(Hierarchical Policer\) | 920](#)

[premium | 1055](#)

alt-name (Application Identification)

Syntax

```
alt-name alt-name
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Provide an alternate name for the application.

Options

alt-name—Alternate name for the application.

Range: 1 through 255 characters

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

analytics

List of Syntax

[Syntax \(EX Series and QFX Series\) on page 720](#)

[Syntax \(MX Series & PTX Series\) on page 724](#)

Syntax (EX Series and QFX Series)

Junos OS Release 13.2X51-D15 and later:

```
analytics {
  collector {
    local {
      file filename {
        size size;
        files number;
      }
    }
  }
  address ip-address {
    port number {
      transport protocol {
        export-profile profile-name;
      }
    }
  }
}
export-profiles {
  profile-name {
    interface {
      information;
      statistics {
        queue;
        traffic;
      }
      status {
        link;
        queue;
        traffic;
      }
    }
  }
  stream-format format;
  system {
    information;
    status {
      queue;
```

```

        traffic;
    }
}
}
}
resource {
    interfaces {
        interface-name {
            resource-profile name;
        }
    }
    system {
        polling-interval {
            queue-monitoring interval;
            traffic-monitoring interval;
        }
        resource-profile name;
    }
}
resource-profiles {
    profile-name {
        depth-threshold {
            high number;
            low number;
        }
        latency-threshold {
            high number;
            low number;
        }
        no-queue-monitoring;
        no-traffic-monitoring;
        queue-monitoring;
        traffic-monitoring;
    }
}
traceoptions {
    file filename {
        files number;
        size size;
    }
}
}

```


Junos OS Release 13.2X50-D15 and 13.2X51-D10 only:

```

analytics {
  interfaces {
    all {
      depth-threshold high number low number;
      latency-threshold high number low number;
      queue-statistics;
      no-queue-statistics;
      traffic-statistics;
      no-traffic-statistics;
    }
    interface-name {
      depth-threshold high number low number;
      latency-threshold high number low number;
      queue-statistics;
      no-queue-statistics;
      traffic-statistics;
      no-traffic-statistics;
    }
  }
  queue-statistics {
    file filename {
      files number-of-files;
      size size;
    }
    interval interval;
  }
  streaming-servers {
    address ip-address {
      port number {
        stream-format format;
        stream-type type
      }
    }
  }
  traceoptions {
    file filename {
      files number;
      size size;
    }
  }
  traffic-statistics {
    file filename {
      files number-of-files;

```

```
        size size;  
    }  
    interval interval;  
}  
}
```

Syntax (MX Series & PTX Series)

```

analytics {
  agent {
    service-agents {
      agent-name {
        inputs {
          analytics {
            parameters {
              generate-tags value;
              sample-frequency value;
              sensors file-path;
            }
          }
          input-ipfix {
            parameters {
              maximum-connections number;
              tcp-port port-number;
              vrf-name name;
            }
          }
          input-jti-ipfix {
            parameters {
              record-group group-name {
                record ipfix-record-name;
                reporting-interval seconds;
              }
            }
          }
        }
      }
      outputs {
        file {
          parameters {
            path file-path;
          }
        }
        kafka {
          parameters {
            server ip-address;
            topic topic-name;
            encoding encoding-type;
          }
        }
        output-ipfix {
          parameters {

```


RELATED DOCUMENTATION

Network Analytics Overview

[Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data | 672](#)

Configuring NTF Agent

[IPFIX Mediation on the BNG | 667](#)

[Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator | 680](#)

[Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector | 677](#)

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	183
Configuring CoS Adjustment Control Profiles	186
Verifying the CoS Adjustment Control Profile Configuration	188
adjustment-control-profiles	712
application (Adjustment Control Profiles)	729

application (Adjustment Control Profiles)

Syntax

```
application {  
  (anyp | dhcp-tags | pppoe-tags | radius-coa)  
  priority priority;  
  algorithm algorithm;  
}
```

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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](#) | 183

[Configuring CoS Adjustment Control Profiles](#) | 186

[Verifying the CoS Adjustment Control Profile Configuration](#) | 188

application (Application Identification)

Syntax

```

application application-name <description description> {
  address-mapping name {
    destination {
      ip ip-address-prefix;
    }
    source {
      ip ip-address-prefix;
    }
    order order;
    order-priority (high | low);
  }
}
alt-name alt-name;
cacheable;
compatibility junos-compatibility-version;
description description;
icmp-mapping {
  code icmp-code;
  order order;
  order-priority (high | low);
  type icmp-type;
}
ip-protocol-mapping {
  order order;
  order-priority (high | low);
  protocol protocol-number
}
order order;
over protocol-type {
  signature I4-I7-signature-name {
    chain-order
    member member-name {
      check-bytes max-bytes-to-check;
      context context;
      pattern pattern;
      direction direction;
    }
    order order;
    order-priority (high | low);
    port-range {
      tcp [port-range];

```

```

        udp [port-range];
    }
    protocol (http | ssl | tcp | udp);
]
priority;
tags tag-value;
type type;
}

```

Hierarchy Level

```
[edit services application-identification]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure identification of an application for which one or more custom signatures are defined.

Options

application-name—Name of the application for which one or more custom signatures has been defined.

description—(Optional) Textual description of the application for which mappings are provided.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview](#) | 431

[Configuring Custom Application Signatures](#) | 435

application-groups (PCC Rules)

Syntax

```
application-groups [application-group-name];
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rule-name from],  
[edit services pcef pcc-rules rule-name from]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-rules *rule-name* from]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify one or more application groups to define the match criteria for the policy and charging control (PCC) rule. You can specify a maximum of 10 application groups in a PCC rule.

NOTE: You must also include the **flows** statement. If you do not want to filter subscriber traffic based on service data flow filters, use **flows any**.

If you are using Junos OS Subscriber Aware, specify the name of the application group at the **[edit unified-edge pcef pcc-rules *rule-name* from]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the application group at the **[edit services pcef pcc-rules *rule-name* from]** hierarchy level.

Options

application-group-name—Name of an application group that is used to detect IP packet flows.

Range: 1 through 63 characters.

NOTE: The referenced application groups must have been previously configured in the **[edit services application-identification]** hierarchy level.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Policy and Charging Control Rules](#) | 418

application-identification (Application Identification)

Syntax

```

application-identification {
  application application-name <description description> {
    address-mapping name {
      destination {
        ip ip-address-prefix;
      }
      source {
        ip ip-address-prefix;
      }
      order order;
      order-priority (high | low);
    }
  }
  alt-name alt-name;
  cacheable;
  compatibility junos-compatibility-version;
  description description;
  icmp-mapping {
    code icmp-code;
    order order;
    order-priority (high | low);
    type icmp-type;
  }
  ip-protocol-mapping {
    order order;
    order-priority (high | low);
    protocol protocol-number
  }
  order
  over protocol-type {
    signature I4-I7-signature-name {
      chain-order
      member member-name {
        check-bytes max-bytes-to-check;
        context context;
        pattern pattern;
        direction direction;
      }
      order order;
      order-priority (high | low);
      port-range {

```

```

        tcp [port-range];
        udp [port-range];
    }
    protocol (http | ssl | tcp | udp);
]
}
priority;
tags tag-value;
type type;
}
download {
    ignore-server-validation;
}
enable-performance-mode;
}

```

Hierarchy Level

[edit services]

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure application identification options to identify the application as it passes through the device.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

application-identification-profile (Service Set)

Syntax

```
application-identification-profile app-id-profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Specify the dummy application identification profile that you configured at the **[edit services application-identification-profile]** hierarchy level. This profile is a placeholder profile with no configuration options, but it must be specified to enable application identification functionality on the services plane.

Options

app-id-profile-name—Name of the application identification profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Applying Services to Subscriber-Aware Traffic with a Service Set

[Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control](#) | 424

applications (PCC Rules)

Syntax

```
applications [application-name];
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rule-name from],  
[edit services pcef pcc-rules rule-name from]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-rules *rule-name* from]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify one or more applications to define the match criteria for the policy and charging control (PCC) rule. You can specify a maximum of 10 applications in a PCC rule.

NOTE: You must also include the **flows** statement. If you do not want to filter subscriber traffic based on service data flow filters, use **flows any**.

If you are using Junos OS Subscriber Aware, specify the name of the applications at the **[edit unified-edge pcef pcc-rules *rule-name* from]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the applications at the **[edit services pcef pcc-rules *rule-name* from]** hierarchy level.

Options

application-name—Name of one or more applications that is used to detect IP packet flows.

Range: 1 through 63 characters.

NOTE: The referenced application must have been previously configured in the **[edit services application-identification]** hierarchy level.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Policy and Charging Control Rules](#) | 418

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

authentication-order

Syntax

```
authentication-order [ authentication-methods ];
```

Hierarchy Level

```
[edit access profile profile-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

none option added in Junos OS Release 11.2.

nasreq option added in Junos OS Release 16.1.

s6a option added in Junos OS Release 19.3R1.

Description

Set the order in which AAA tries different authentication methods when verifying that a client can access the router or switch. For each login attempt, AAA tries the authentication methods in order, from first to last.

A given subscriber does not undergo both authentication and authorization as separate steps. When both **authentication-order** and **authorization-order** are specified, DHCP subscribers honor the configured authorization order, all other subscribers use the configured authentication-order.

Starting in Junos OS Release 18.2R1, the **password** option can also be used to specify that local authentication and local authorization is attempted for individual subscribers that are configured with the **subscriber** statement at the `[edit access profile profile-name]` hierarchy level.

Options

authentication-methods—Ordered list of methods to use for authentication attempts. The list includes one or more of the following methods in any combination:

- **nasreq**—Verify subscribers using the Diameter-based Network Access Server Requirements (NASREQ) protocol.
- **none**—No authentication is performed. Grants authentication without examining the client credentials. Can be used, for example, when the Diameter function Gx-Plus is employed for notification during subscriber provisioning.

NOTE: Subscriber access management does not support the **none** option; authentication fails when this option is specified.

- **password**—Verify the client using the information configured at the [edit access profile *profile-name* client *client-name*] hierarchy level.

Subscriber access management does not support the **password** option until Junos OS Release 18.2R1. Starting in Junos OS Release 18.2R1, this option is used to enable local authentication and optionally local authorization for individual subscribers. Local authentication is typically used when you do not have external authentication and authorization servers. The password itself must be configured with the **subscriber** statement in the same access profile. Local authentication is performed when a subscriber logs in with a matching username; it succeeds if the subscribers login password matches the password in the profile.

If you have external authentication and authorization servers, you can use local authentication as a backup authentication method. In this case, configure **password** other than first in the list of methods.

- **radius**—Verify the client using RADIUS authentication services.
- **s6a**—Verify subscribers using the Diameter-based s6a protocol.

Default: password

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring CHAP Authentication with RADIUS

Specifying the Authentication and Accounting Methods for Subscriber Access

Configuring Access Profiles for L2TP or PPP Parameters

Configuring Local Authentication and Authorization for Subscribers

Configuring S6a

bandwidth (Tunnel Services)

Syntax

```
bandwidth bandwidth-value;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number tunnel-services]
```

Release Information

Statement introduced in Junos OS Release 8.2.

Statement introduced in Junos OS Release 12.3X54 for ACX Series routers.

Description

(ACX Series, MX Series 5G Universal Routing Platforms and T4000 Core Routers only) Configure the amount of bandwidth in gigabits per second reserved on each Packet Forwarding Engine for tunnel traffic using tunnel services. Configuring the bandwidth creates a virtual tunnel interface that is represented as **lt-*<fpc/pic/port>***.

Options

bandwidth-value—Amount of bandwidth in Gbps to reserve for tunnel traffic using tunnel services:

- On ACX Series routers, the bandwidth values can be **1g** or **10g**.
- On MX Series routers, the bandwidth values can be as follows:
 - **1g**
 - **10g** through **100g** in 10 Gbps increments: **10g, 20g, 30g, 40g, 50g, 60g, 70g, 80g, 90g, 100g**
 - **100g** through **400g** in 100 Gbps increments: **100g, 200g, 300g, 400g**
- On T4000 routers, the bandwidth values can be **10g** through **100g** in 10 Gbps increments: **10g, 20g, 30g, 40g, 50g, 60g, 70g, 80g, 90g, 100g**.

NOTE: The bandwidth that you specify determines the port number of the tunnel interfaces that are created. When you specify a bandwidth of **1g**, the port number is always 10. When you specify any other bandwidth, the port number is always 0.

NOTE: If you specify a bandwidth that is not compatible with the type of DPCs or MPCs and their respective Packet Forwarding Engine, tunnel services are not activated. For example, you cannot specify 1 gigabit per second bandwidth for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC or 16x10GE 3D MPC.

When the tunnel bandwidth is unspecified in the Routing Engine CLI, the maximum tunnel bandwidth for MPC3E is 60G.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Tunnel Interfaces on a Gigabit Ethernet 40-Port DPC

Tunnel Interface Configuration on MX Series Routers Overview

Configuring Tunnel Interfaces on T4000 Routers

Example: Configuring Tunnel Interfaces on a 10-Gigabit Ethernet 4-Port DPC

Example: Configuring Tunnel Interfaces on the MPC3E

[tunnel-services \(Chassis\) | 1227](#)

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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 and T Series routers) 8000 through 100,000,000,000
- (Mx Series routers) 8000 through 18,446,744,073,709,551,615

NOTE: When you specify a numeric value beyond the supported bandwidth of the PFE, the router caps the bandwidth at the maximum supported bandwidth of the PFE.

Default: None.

Required Privilege Level

- firewall—To view this statement in the configuration.
- firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Two-Color Policer Configuration Overview</i>
<i>Policer Bandwidth and Burst-Size Limits</i>
<i>Policer Color-Marking and Actions</i>
<i>Single Token Bucket Algorithm</i>
<i>Determining Proper Burst Size for Traffic Policers</i>
bandwidth-percent 747
burst-size-limit (Policer) 754

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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 priority. You apply a hierarchical policer to ingress Layer 2 traffic to allow 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: The bandwidth percentage policer cannot be used to rate-limit tunnel or software interfaces, or for forwarding table filters. It is only valid for interface-specific filters. When used for matching bandwidth or burst-size on aggregated Ethernet or SONET bundles, bandwidth percentage policers must be used in conjunction with **shared-bandwidth-policer**.

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

<i>Two-Color Policer Configuration Overview</i>
<i>Policer Bandwidth and Burst-Size Limits</i>
<i>Policer Color-Marking and Actions</i>
<i>Single Token Bucket Algorithm</i>
<i>Determining Proper Burst Size for Traffic Policers</i>
<i>Bandwidth Policers</i>
bandwidth-limit (Policer) 745
burst-size-limit (Policer) 754

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](#) | 40

[Configuring Schedulers in a Dynamic Profile for Subscriber Access](#) | 52

[scheduler \(Dynamic Scheduler Maps\)](#) | 1134

burst-size-limit (Hierarchical Policer)

Syntax

```
burst-size-limit bytes;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer aggregate if-exceeding],
[edit dynamic-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

On M40e, M120, and M320 (with FFPC and SFPC) edge routers; on MPCs hosted on MX Series routers; on T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs; and on 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 (1500 through 100,000,000,000 on MPCs hosted on MX Series routers)

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 \(Hierarchical Policer\)](#) | **718**

bandwidth-limit (Hierarchical Policer)

[premium \(Hierarchical Policer\)](#) | **1055**

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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 allow 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.

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 and EX Series switches, 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 and EX Series switches. For a single-rate two-color policer on an MX Series router and on an EX Series switch, 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.

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 \(Policer\) | 745](#)

[bandwidth-percent | 747](#)

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](#) | 183

[Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates](#) | 111

[Bandwidth Management for Downstream Traffic in Edge Networks Overview](#) | 109

[egress-shaping-overhead](#)

cacheable (Application Identification)

Syntax

```
cacheable
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Enable the application system cache (ASC), which saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service. The ASC is disabled by default.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

captive-portal-content-delivery (Captive Portal Content Delivery)

Syntax

```

captive-portal-content-delivery {
  auto-deactivate value;
  profile name
    cpcd-rule-sets rule-set-name;
    cpcd-rules rule-name;
    dynamic;
    http-redirect-options url;
    ipda-rewrite-options {
      destination-address destination-address;
      destination-port destination-port;
    }
  }
  rule rule-name {
    match-direction (input | output | input-output);
    term term-name {
      from {
        destination-address address <except>;
      }
      then {
        accept;
        insert tag tag-name tag-value tag-value;
        redirect url;
        rewrite destination-address address <destination-port port-number>;
        syslog;
      }
    }
  }
  rule-set rule-set-name {
    rule rule-name;
  }
  traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-readable)>;
    flag name;
    no-remote-trace no-remote-trace;
  }
}

```

Hierarchy Level

[edit dynamic-profiles *profile-name* **services**],

[edit [services](#)]

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the [edit dynamic-profiles *profile-name* services] hierarchy level added in Junos OS Release 17.2R1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

auto-deactivate option added in Junos OS Release 19.4R1.

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. Use the statement at the [edit services...] hierarchy level for static CPCD. Use the statement at the [edit dynamic-profiles *profile-name* services...] hierarchy level for converged services CPCD.

The **profile**, **rule-set**, and **traceoptions** stanzas are not supported at the [edit dynamic-profiles *profile-name* hierarchy level].

Options

auto-deactivate value—(Optional) Specify one of the following values to determine whether the redirect service is removed automatically when the router receives the subscriber's initial HTTP-GET message:

- **user-defined-variable**—To use this option in a dynamic profile, you must create a user-defined variable with a name of your choice. The value of the variable can be supplied by the RADIUS server or PCRF. You can also define a default value that is used when the external servers do not supply it. Use the *variables* statement in the dynamic profile to define the default value. Whether supplied by the external server or by the CLI, the value must be either **initial-get** or **never**.
- **initial-get**—Receipt of the subscriber's initial HTTP-GET message triggers removal of the redirect service.
- **never**—Removal of the redirect service is triggered only by the external server, such as by a CoA message from a RADIUS server.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview | 473](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

[Adding Subscriber Information to HTTP Redirect URLs | 529](#)

[How to Automatically Remove the HTTP Redirect Service After the Initial Redirect | 532](#)

captive-portal-content-delivery-profile (Services)

Syntax

```
captive-portal-content-delivery-profile profile-name
  interface-service {
    service-interface name;
  }
  next-hop-service {
    inside-service-interface interface-name.unit-number;
    outside-service-interface interface-name.unit-number;
    outside-service-interface-type local;
  }
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced before Junos OS Release 17.2.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

Description

NOTE: Starting in Junos OS Release 17.2R1, you can configure converged services for MS-MPCs and MS-MICs. Starting in Junos OS Release 19.3R2, you can configure converged services for SPC3s if you have enabled Next Gen Services on the MX Series router. At the **edit service-set service set name captive-portal-content-delivery-profile profile-name interface-service** hierarchy level, you can configure captive portal content delivery (CPCD) profiles for MS-MICs and MS-MPCs by including the **service-interface ms-fpc/pic/port** statement, and configure captive portal content delivery (CPCD) profiles for SPC3s by including the **service-interface vms-fpc/pic/port** statement.

Options

service-interface *name*—Name of the service device associated with the interface-wide service set.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview](#) | 473

dynamic-profile

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.

RELATED DOCUMENTATION

CoS Adjustment Control Profiles Overview 183
Configuring CoS Adjustment Control Profiles 186
adjustment-control-profiles 712
Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates 111
Bandwidth Management for Downstream Traffic in Edge Networks Overview 109
egress-shaping-overhead
bytes 756
frame-mode 892

chain-order (Application Identification)

Syntax

```
chain-order;
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type signature I4-I7-signature-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Read members in order. By default, chain ordering is turned off. If there is only one member, this option is ignored.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

check-bytes (Application Identification)

Syntax

```
check-bytes max-bytes-to-check;
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type signature l4-l7-signature-name  
  member member-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Specify the maximum number of bytes to be inspected. This statement applies to TCP and UDP protocols for stream context. It is not considered for other protocols and contexts.

Options

max-bytes-to-check—Number of bytes to be inspected.

Range: 1 through 5000

Default: Not configured

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

class (Defining Login Classes)

Syntax

```
class class-name {
    access-end hh:mm;
    access-start hh:mm;
    ( allow-commands "(regular-expression1)|(regular-expression2)..." | allow-commands-regexps ["regular expression
        1" "regular expression 2 " ... ] );
    ( allow-configuration "(regular-expression1)|(regular-expression2)..." | allow-configuration-regexps ["regular expression
        1" "regular expression 2 " ... ] );
    allow-hidden-commands;
    allow-sources [ source-addresses ... ];
    allow-times [ times ... ];
    allowed-days [ days of the week ];
    cli {
        prompt prompt;
    }
    configuration-breadcrumbs;
    confirm-commands ["regular expression or command 1" "regular expression or command 2" ...] {
        confirmation-message;
    }
    ( deny-commands "(regular-expression1)|(regular-expression2)..." | deny-commands-regexps ["regular expression
        1" "regular expression 2 " ... ] );
    ( deny-configuration "(regular-expression1)|(regular-expression2)..." | deny-configuration-regexps ["regular expression
        1" "regular expression 2 " ... ] );
    deny-sources [ source-addresses ... ];
    deny-times [ times ... ];
    idle-timeout minutes;
    logical-system logical-system-name;
    login-alarms;
    login-script login-script;
    login-tip;
    no-hidden-commands {
        except ["regular expression or command 1" "regular expression or command 2" ...];
    }
    no-scp-server;
    no-sftp-server;
    permissions [ permissions ];
    satellite all;
    security-role (audit-administrator | crypto-administrator | ids-administrator | security-administrator);
    tenant tenant-system-name;
}
```

Hierarchy Level

[edit system [login](#)]

Release Information

The **class**, **allow-commands**, **deny-commands**, **allow-configuration**, **deny-configuration**, **idle-timeout**, **login-alarms**, **login-tip**, and **permissions** statements were introduced before Junos OS Release 7.4.

All of the previously mentioned statements were introduced in Junos OS Release 9.0 for the EX Series.

The **login-script** statement was introduced in Junos OS Release 9.5.

The **access-end**, **access-start**, and **allowed-days** statements were introduced in Junos OS Release 10.1.

All of the previously mentioned statements were introduced in Junos OS Release 11.1 for the QFX Series.

All of the previously mentioned statements were introduced in Junos OS Release 11.2 for the SRX Series.

The **allow-configuration-regexps**, **deny-configuration-regexps**, and **security-role** statements were introduced in Junos OS Release 11.2.

The **configuration-breadcrumbs** statement was introduced in Junos OS Release 12.2.

All of the previously mentioned statements were introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

All of the previously mentioned statements were introduced in Junos OS Release 15.1X49-D70 for the vSRX, SRX4100, SRX4200 and SRX1500 devices.

All of the previously mentioned statements were introduced in Junos OS Release 16.1 for the MX Series and PTX Series.

The **allow-hidden-commands**, **confirm-commands**, **no-hidden-commands**, and **satellite** statements were introduced in Junos OS Release 16.1.

The **cli** statement was introduced in Junos OS Release 17.3.

The **allow-commands-regexps** and **deny-commands-regexps** statements were introduced in Junos OS Release 18.1.

The **tenant** statement was introduced in Junos OS 18.4.

The **no-scp-server** and **no-sftp-server** statements were introduced in Junos OS Release 19.2.

Description

Define a login class. All users who log in to the router or switch must be in a login class. Therefore, you must define a Junos OS login class for each user or type of user. You can define any number of login classes depending on the types of permissions the users need. You may not need to define any login classes; Junos OS has several predefined login classes, to suit a variety of needs. However, the predefined login classes cannot be modified. If you define a class with the same name as a predefined class, Junos OS appends **-local** to the login class name and creates a new login class. See *Predefined System Login Classes* for more information.

Options

class-name— A name you choose for the login class.

access-end— Specify the end time in **HH:MM** (24-hour) format, where **HH** represents the hours and **MM** represents the minutes.

NOTE: Access start and end times that span across 12:00 AM starting on a specified day results in the user having access until the next day, even if the access day is not explicitly configured on the **allowed-days** statement.

access-start— Specify the start time in **HH:MM** (24-hour) format, where **HH** represents the hours and **MM** represents the minutes.

NOTE: Access start and end times that span across 12:00 AM starting on a specified day results in the user having access until the next day, even if the access day is not explicitly configured on the **allowed-days** statement.

(**allow-commands** | **allow-commands-regexps**)— Specify one or more regular expressions to allow users in this class to issue operational mode commands. You use the **allow-commands** or the **allow-commands-regexps** statement to explicitly allow authorization for commands that would otherwise be denied by the access privilege levels for a login class.

For the **allow-commands** statement, each expression separated by a pipe (|) symbol must be a complete standalone expression, and must be enclosed in parentheses (). Do not use spaces between regular expressions separated with parentheses and connected with the pipe (|) symbol.

For the **allow-commands-regexps** statement, you configure a set of strings in which each string is a regular expression, enclosed in double quotes and separated with a space operator. Each string is evaluated against the full path of the command, which provides faster matching than the **allow-command** statement. You can also include values for variables in the regular expressions, which is not supported using the **allow-commands** statement.

The **deny-commands** or the **deny-commands-regexps** statement takes precedence if it is used in the same login class definition.

NOTE: The **allow/deny-commands** and **allow/deny-commands-regexps** statements are mutually exclusive and cannot be configured together for a login class. At a given point in time, a login class can include either the **allow/deny-commands** statements, or the **allow/deny-commands-regexps** statements. If you have existing configurations using the **allow/deny-commands** statements, using the same configuration options with the **allow/deny-commands-regexps** statements might not produce the same results, as the search and match methods differ in the two forms of these statements.

Authorizations can also be configured remotely by specifying Juniper Networks vendor-specific TACACS+ attributes in your authentication server's configuration. For a remote user, when the authorization parameters are configured both remotely and locally, authorization parameters configured remotely and locally are both considered together for authorization. For a local user, only the authorization parameters configured locally for the class are considered.

Default: If you do not configure authorizations for operational mode commands using the **allow/deny-commands** or **allow/deny-commands-regexps** statements, users can edit only those commands for which they have access privileges set with the **permissions** statement.

Syntax: *regular-expression*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Enter as many expressions as needed.

(**allow-configuration** | **allow-configuration-regexps**)— Specify one or more regular expressions to explicitly allow users in this class to access the specified levels in the configuration hierarchy even if the permissions set with the **permissions** statement do not grant such access.

For the **allow-configuration** statement, each expression separated by a pipe (|) symbol must be a complete standalone expression, and must be enclosed in parentheses (). Do not use spaces between regular expressions separated with parentheses and connected with the pipe (|) symbol.

For the **allow-configuration-regexps** statement, you configure a set of strings in which each string is a regular expression, enclosed in double quotes and separated with a space operator. Each string is evaluated against the full path of the command, which provides faster matching than the **allow/deny-configuration** statements. You can also include values for variables in the regular expressions, which is not supported using the **allow/deny-configuration** statements.

The **deny-configuration** or **deny-configuration-regexps** statement takes precedence if it is used in the same login class definition.

NOTE: The **allow/deny-configuration** and **allow/deny-configuration-regexps** statements are mutually exclusive and cannot be configured together for a login class. At a given point in time, a login class can include either the **allow/deny-configuration** statements, or the **allow/deny-configuration-regexps** statements. If you have existing configurations using the **allow/deny-configuration** statements, using the same configuration options with the **allow/deny-configuration-regexps** statements might not produce the same results, as the search and match methods differ in the two forms of these statements.

Default: If you omit the **allow-configuration/allow-configuration-regexps** statement and the **deny-configuration/deny-configuration-regexps** statement, users can edit only those commands for which they have access privileges through the **permissions** statement.

Syntax: *regular-expression*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Enter as many expressions as needed.

allow-hidden-commands— Allow all hidden commands to be run. If the **no-hidden-commands** statement is specified at the [edit system] hierarchy level, overrides that restriction for this login class. Hidden commands are Junos OS commands that are not published but could be run on a router. Hidden commands serve a specific purpose, but for most part are not expected to be used, and as such are not actively supported. The **no-hidden-commands** statement at the [edit system] hierarchy level allows you to block all hidden commands to all users except the root users.

Default: Hidden commands are enabled by default.

allow-sources [*source-addresses ...*]— Restrict incoming remote access to only particular hosts. Specify one or more source addresses from which access is allowed. The source addresses can be IPv4 or IPv6 addresses, prefix lengths, or hostnames.

allow-times [*times...*]— Restrict remote access to certain times.

allowed-days [*days of the week*]— Specify one or more days of the week when users in this class are allowed to log in.

Values:

- monday—Monday
- tuesday—Tuesday
- wednesday—Wednesday
- thursday—Thursday
- friday—Friday
- saturday—Saturday
- sunday—Sunday

cli— Set the CLI prompt specified for the login class. If a CLI prompt is also set at the [edit system login user cli] hierarchy level, the prompt set for the login user has precedence over the prompt set for the login class.

prompt *prompt*— Specify the prompt string you want to see displayed in the CLI prompt.

configuration-breadcrumbs— Enable the configuration breadcrumbs view in the CLI to display the location in the configuration hierarchy. For an example of how to enable this view, see *Example: Enabling Configuration Breadcrumbs* .

confirm-commands— Specify that confirmation for particular commands is explicitly required and, optionally, specify the wording of the message displayed at confirm time. You can specify the commands using a list of regular expressions or commands.

Syntax: *message*

Default: If you omit this option, then confirmation for commands is not required. If the optional message is not set, then the default "Do you want to continue?" message is displayed.

(**deny-commands** | **deny-commands-regexps**)— Specify one or more regular expressions to explicitly deny users in this class permission to issue operational mode commands, even though the permissions set with the **permissions** statement would allow it.

For the **deny-commands** statement, each expression separated by a pipe (|) symbol must be a complete standalone expression, and must be enclosed in parentheses (). Do not use spaces between regular expressions separated with parentheses and connected with the pipe (|) symbol.

For the **deny-commands-regexps** statement, you configure a set of strings in which each string is a regular expression, enclosed in double quotes and separated with a space operator. Each string is evaluated against the full path of the command, which provides faster matching than the **allow/deny-command** statements. You can also include values for variables in the regular expressions, which is not supported using the **allow/deny-commands** statements.

Expressions configured with the **deny-commands** or the **deny-commands-regexps** statement take precedence over expressions configured with **allow-commands/allow-commands-regexps** if the two statements are used in the same login class definition.

NOTE: The **allow/deny-commands** and **allow/deny-commands-regexps** statements are mutually exclusive and cannot be configured together for a login class. At a given point in time, a login class can include either the **allow/deny-commands** statements, or the **allow/deny-commands-regexps** statements. If you have existing configurations using the **allow/deny-commands** statements, using the same configuration options with the **allow/deny-commands-regexps** statements might not produce the same results, as the search and match methods differ in the two forms of these statements.

Authorizations can also be configured remotely by specifying Juniper Networks vendor-specific TACACS+ attributes in your authentication server's configuration. For a remote user, when the authorization parameters are configured both remotely and locally, authorization parameters configured remotely and locally are both considered together for authorization. For a local user, only the authorization parameters configured locally for the class are considered.

Default: If you do not configure authorizations for operational mode commands using **allow/deny-commands** or **allow/deny-commands-regexps**, users can edit only those commands for which they have access privileges set with the **permissions** statement.

Syntax: *regular-expression*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Enter as many expressions as needed.

(**deny-configuration** | **deny-configuration-regexps**)— Specify one or more regular expressions to explicitly deny users in this class access to the specified levels in the configuration hierarchy even if the permissions set with the **permissions** statement grant such access. Note that the user cannot view a particular hierarchy if configuration access is denied for that hierarchy.

For the **deny-configuration** statement, each expression separated by a pipe (|) symbol must be a complete standalone expression, and must be enclosed in parentheses (). Do not use spaces between regular expressions separated with parentheses and connected with the pipe (|) symbol.

For the **deny-configuration-regexps** statement, you configure a set of strings in which each string is a regular expression, enclosed in double quotes and separated with a space operator. Each string is evaluated against the full path of the command, which provides faster matching than the **allow/deny-configuration** statements. You can also include values for variables in the regular expressions, which is not supported using the **allow/deny-configuration** statements.

Expressions configured with **deny-configuration/deny-configuration-regexps** take precedence over expressions configured with **allow-configuration/allow-configuration-regexps** if the two statements are used in the same login class definition.

NOTE: The **allow/deny-configuration** and **allow/deny-configuration-regexps** statements are mutually exclusive and cannot be configured together for a login class. At a given point in time, a login class can include either the **allow/deny-configuration** statements, or the **allow/deny-configuration-regexps** statements. If you have existing configurations using the **allow/deny-configuration** statements, using the same configuration options with the **allow/deny-configuration-regexps** statements might not produce the same results, as the search and match methods differ in the two forms of these statements.

Default: If you omit the **deny-configuration/deny-configuration-regexps** statement and the **allow-configuration/allow-configuration-regexps** statement, users can edit those levels in the configuration hierarchy for which they have access privileges through the **permissions** statement.

Syntax: *regular-expression*—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks. Enter as many expressions as needed.

deny-sources [*source-addresses*]
Never allow remote access from these hosts. The source addresses can be IPv4 or IPv6 addresses, prefix lengths, or hostnames.

deny-times [*times*]
Never allow remote access during these times.

idle-timeout
For a login class, configure the maximum time in minutes that a session can be idle before the session times out and the user is logged out of the device. The session times out after remaining at the CLI operational mode prompt for the specified time.

NOTE: After the user logs in to a device from a shell prompt such as `csh`, if the user starts another program to run in the foreground of the CLI, the idle-timer control is stopped from being computed. The calculation of the idle time of the CLI session is restarted only after the foreground process exits and the control is returned to the shell prompt. When the restart of the idle-timer control occurs, if no interaction from the user occurs on the shell, the user is automatically logged out after the time set on this statement.

Default: If you omit this statement, a user is never forced off the system after extended idle times.

Syntax: *minutes*—Maximum time in minutes that a session can be idle before a user is logged out.

Range: Range: 0 through 4294967295 minutes

NOTE: The idle-timeout feature is disabled if the value of *minutes* is set to 0.

login-alarms
Display system alarms when a user with **admin** permissions logs in to the device. For more information about configuring this statement, see *Configuring System Alarms to Appear Automatically Upon Login*.

login-script
Run the specified op script when a user belonging to the class logs in to the CLI. The script must be enabled in the configuration.

logical-system
Assign the users in this login class to a logical system. If you specify a logical system, you can't include the satellite configuration statement in the configuration for this login class.

login-tip
Display CLI tips when logging in.

Default: If this statement is not configured, CLI tips are not displayed.

no-hidden-commands
Deny all hidden commands, except for those specified, for users in this login class. Each command listed as an exception must be enclosed in quotation marks.

Default: Hidden commands are enabled by default.

Syntax: except [*"command 1" "command 2"...*]

no-scp-server— Disable incoming SCP connections for this login class.

no-sftp-server— Disable incoming SFTP connections for this login class.

permissions— Specify login access privileges for the login class.

Syntax: *permissions*—One or more permission flags, which together specify the access privileges for the login class. Permission flags are not cumulative, so for each class, you must list all the permission flags needed, including **view** to display information and **configure** to enter configuration mode. For a list of permission flags, see *Login Class Permission Flags*.

satellite— Specify access to Junos Fusion satellite devices for the login class. All users assigned to the login class are satellite users. If you include this statement, you can't include the logical-system configuration statement in the configuration for this login class.

Values:

- **all**—Specify all Junos Fusion satellite devices.

security-role— Specify one or more Common Criteria (ISO/IEC 15408) security roles for the login class.

Values:

audit-administrator— Specify which users are responsible for the regular review of specific target of evaluation (TOE) audit data and audit trail deletion. Audit administrators can also invoke the non-cryptographic self-test.

crypto-administrator— Specify which users are responsible for the configuration and maintenance of cryptographic elements related to the establishment of secure connections to and from the TOE audit data.

ids-administrator— Specify which users can act as intrusion detection service (IDS) administrators, who are responsible for all of the activities regarding identity and access management of the organization's employees.

security-administrator— Specify which users are responsible for ensuring that the organization's security policy is in place.

tenant— Assign the users in this class to a tenant system. Tenant systems are used when you need to separate departments, organizations, or customers and each of them can be limited to one virtual router. The main difference between a logical system and a tenant system is that a logical system supports advanced routing functionality using multiple routing instances. In comparison, a tenant system supports only one routing instance, but supports the deployment of significantly more tenants per system.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Defining Junos OS Login Classes</i>
<i>Configuring Time-Based User Access</i>
<i>Understanding Junos OS Access Privilege Levels</i>
<i>Example: Configuring User Permissions with Access Privileges for Operational Mode Commands</i>
<i>Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands, Configuration Statements, and Hierarchies</i>
<i>Example: Configuring User Permissions with Access Privileges for Configuration Statements and Hierarchies</i>
<i>Example: Configuring User Permissions with Access Privilege Levels</i>
<i>Configuring System Alarms to Appear Automatically Upon Login</i>
<i>Executing an Op Script on the Local Device</i>
<i>Understanding Administrative Roles</i>
<i>Example: Configuring Administrative Roles</i>
<i>Tenant Systems Overview</i>
user 1243

class-of-service (Dynamic Profiles)

Syntax

```

class-of-service {
  dynamic-class-of-service-options {
    vendor-specific-tags tag;
  }
  interfaces {
    interface-name ;
  }
  unit logical-unit-number {
    classifiers {
      type (classifier-name | default);
    }
    output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
    report-ingress-shaping-rate bps;
    rewrite-rules {
      dscp (rewrite-name | default);
      dscp-ipv6 (rewrite-name | default);
      ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
      inet-precedence (rewrite-name | default);
    }
  }
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  (scheduler-name) {
    buffer-size (seconds | percent percentage | remainder | temporal microseconds);
    drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any | non-tcp | tcp)
      drop-profile profile-name;
    excess-priority (low | high | $junos-cos-scheduler-excess-priority);
    excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
    overhead-accounting (shaping-mode) <bytes (byte-value>;
    priority priority-level;
    shaping-rate (rate | predefined-variable);
    transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
  }
}
traffic-control-profiles profile-name {

```

```

adjust-minimum rate;
delay-buffer-rate (percent percentage | rate);
excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
excess-rate-high (percent percentage | proportion value);
excess-rate-low (percent percentage | proportion value);
guaranteed-rate (percent percentage | rate) <burst-size bytes>;
max-burst-size cells;
overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
peak-rate rate;
scheduler-map map-name;
shaping-rate (percent percentage | rate | predefined-variable) <burst-size bytes>;
shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
sustained-rate rate;
}
}

```

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 client profile or a dynamic service 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 | 40](#)

[Configuring Static Hierarchical Scheduling in a Dynamic Profile](#)

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[Applying a Classifier to a Subscriber Interface in a Dynamic Profile | 221](#)

[classifiers \(Definition\)](#)

code (Application Identification)

Syntax

```
code icmp-code;
```

Hierarchy Level

```
[edit services application-identification application application-name icmp-mapping]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Match the specified ICMP code to create a custom application signature.

Options

value—Numeric value for the ICMP code.

Range: 0 through 254

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

collector (LRF Profile)

Syntax

```
collector collector-name {
  destination {
    address collector-address;
    port collector-port-number;
  }
  source-address source-address;
}
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure a collector that receives logging and reporting data. This collector can be specified in LRF rules.

Options

collector-name—Name for the collector.

Range: Up to 32 characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

[Configuring Logging and Reporting for Junos OS Subscriber Aware](#)

[Configuring Logging and Reporting for Subscriber Management | 459](#)

collector (LRF Rule)

Syntax

```
collector collector-name;
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name then report]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Specify the collector that receives the data if the LRF rule is matched.

Options

collector-name—Name of the collector that receives the data. The referenced collector must already be defined at the **[edit services lrf profile *profile-name*]** 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 LRF Profile for Subscribers | 460](#)

[Configuring Logging and Reporting for Junos OS Subscriber Aware](#)

[Configuring Logging and Reporting for Subscriber Management | 459](#)

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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 rate 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 rate 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](#) | [786](#)

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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 rate 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 rate 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](#) | [784](#)

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

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 790
excess-burst-size 847
peak-burst-size 1030
peak-information-rate 1032

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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 on EX, M, and T Series routers
- 1500 through 18,446,744,073,709,551,615 bps on Mx Series routers

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Three-Color Policer Configuration Overview</i>
<i>Policer Bandwidth and Burst-Size Limits</i>
<i>Policer Color-Marking and Actions</i>
<i>Dual Token Bucket Algorithms</i>
<i>Determining Proper Burst Size for Traffic Policers</i>
committed-burst-size 788
excess-burst-size 847
peak-burst-size 1030
peak-information-rate 1032

compatibility (Application Identification)

Syntax

```
compatibility junos-compatibility-version;
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Specify the Junos OS release for compatibility.

Options

junos-compatibility-version—Name of the Junos OS software release compatibility version, such as 17.1.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

connection-limit

Syntax

```
connection-limit limit;
```

Hierarchy Level

```
[edit system services finger],  
[edit system services ftp],  
[edit system services netconf ssh],  
[edit system services ssh],  
[edit system services telnet],  
[edit system services xnm-clear-text],  
[edit system services xnm-ssl]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description

Configure the maximum number of connections sessions for each type of system services (finger, ftp, ssh, telnet, xnm-clear-text, or xnm-ssl) per protocol (either IPv6 or IPv4).

Options

limit—(Optional) Maximum number of established connections per protocol (either IPv6 or IPv4).

Range: 1 through 250

Default: 75

NOTE: The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured **connection-limit** value if the system resources are limited.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring clear-text or SSL Service for Junos XML Protocol Client Applications

Configuring DTCP-over-SSH Service for the Flow-Tap Application

Configuring Finger Service for Remote Access to the Router

Configuring FTP Service for Remote Access to the Router or Switch

Configuring SSH Service for Remote Access to the Router or Switch

Configuring Telnet Service for Remote Access to a Router or Switch

context (Application Identification)

Syntax

```
context context;
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type signature I4-I7-signature-name
member member-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Define a predefined service-specific context as an additional matching criterion for application identification.

Options

context—One of the following predefined contexts:

NOTE: If the MX Series router is running Next Gen Services, then the following restrictions apply:

- Only the **http-header** context types are available at the **[edit services application-identification application *application-name* over http signature *I4-I7-signature-name* member *member-name*]** hierarchy level.
- Only the **ssl-server** context type is available at the **[edit services application-identification application *application-name* over ssl signature *I4-I7-signature-name* member *member-name*]** hierarchy level.
- Only the **stream** context type is available at the **[edit services application-identification application *application-name* over (tcp | udp) signature *I4-I7-signature-name* member *member-name*]** hierarchy level.

- **http-get-url-parsed-param-parsed**—Decoded, normalized GET URL in an HTTP request and the decoded CGI parameters, if any.
- **http-header-content-type**—Content-Type header in an HTTP transaction.
- **http-header-cookie**—Cookie header in an HTTP transaction.
- **http-header-host**—Host header in an HTTP request.

- **http-header-user-agent**—User-agent header in an HTTP transaction.
- **http-post-url-parsed-param-parsed**—Decoded, normalized POST URL in an HTTP request and the decoded CGI parameters, if any.
- **http-post-variable-parsed**—Decoded POST URL or form data variables.
- **http-url-parsed**—Decoded, normalized URL in an HTTP request.
- **http-url-parsed-param-parsed**—Decoded, normalized URL in an HTTP request and the decoded CGI parameters, if any.
- **ssl-server-name**—Server name in the TLS server name extension or in the SSL server certificate.
- **stream**— TCP or UDP stream data.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

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 6,400,000,000,000 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](#) | 40

[Configuring Traffic Scheduling and Shaping for Subscriber Access](#) | 47

[output-traffic-control-profile](#) | 1006

description (Application Identification)

Syntax

```
description description
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Provide a description of the application.

Options

description—Textual description of the application.

Range: 1 through 255 characters

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

destination (Application Identification)

Syntax

```
destination ip ip-address-prefix;
```

Hierarchy Level

```
[edit services application-identification application application-name address-mapping]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Specify the destination IP address for address mapping-based application identification.

Options

ip-address-prefix—IP address and prefix for matching.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

destination (LRF Profile)

Syntax

```
destination {  
  address collector-address;  
  port collector-port-number;  
}
```

Hierarchy Level

```
[edit services lrf profile profile-name collector collector-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Specify the destination IP address and port number of the collector.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers](#) | 460

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management](#) | 459

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

ddos-protection (DDoS)

List of Syntax

[Syntax \(PTX Series Routers and QFX Series Switches\) on page 803](#)

[Syntax \(Other Routers and EX9200 Switches\) on page 803](#)

Syntax (PTX Series Routers and QFX Series Switches)

```
ddos-protection
  global {
    disable-fpc;
    disable-logging;
  }
  protocols protocol-group (aggregate | packet-type) {
    bandwidth packets-per-second;
    burst size;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    fpc slot-number {
      bandwidth-scale percentage;
      burst-scale percentage;
      disable-fpc;
    }
    priority level;
  }
  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;
  }
}
```

Syntax (Other Routers and EX9200 Switches)

```
ddos-protection
  global {
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection;
    flow-level-control;
```

```

    flow-detection-mode;
    flow-report-rate;
    violation-report-rate;
}
protocols protocol-group (aggregate | packet-type) {
    bandwidth packets-per-second;
    burst size;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode (automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
        logical-interface flow-bandwidth;
        physical-interface flow-bandwidth;
        subscriber flow-bandwidth;
    }
    flow-level-control {
        logical-interface flow-control-mode;
        physical-interface flow-control-mode;
        subscriber flow-control-mode;
    }
    flow-level-detection {
        logical-interface flow-detection-mode;
        physical-interface flow-detection-mode;
        subscriber flow-detection-mode;
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc slot-number {
        bandwidth-scale percentage;
        burst-scale percentage;
        disable-fpc;
    }
    no-flow-logging
    priority level;
    recover-time seconds;
    timeout-active-flows;
}

```

```
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]

Release Information

Statement introduced in Junos OS Release 11.2.

Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.

Statement introduced in Junos OS Release 14.1X53 on QFX Series switches.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

Statement introduced in Junos OS Release 17.4R1 on PTX Series routers.

Description

Configure DDoS protection policers for control plane DDoS protection.

DDoS attacks typically use network control packets to trigger large numbers of exceptions to a device's control plane that disrupts normal network operations. DDoS protection polices traffic to enable the device to continue functioning under a DDoS attack.

DDoS protection is enabled by default on supporting devices for the protocol groups and packet types available on the device. You can disable particular policers or change default policer parameters, including:

- Set the maximum allowed traffic rate, maximum burst size, and traffic priority.
- Define how much time must pass since the last violation before the traffic flow is considered to have recovered from the attack.
- Scale bandwidth and burst values for individual line cards so that the policers at this level trigger at lower thresholds than the overall protocol or packet thresholds.

DDoS protection supports policers for many protocol groups and specific packet types within some protocol groups. Protocol group and packet type support varies across platforms and Junos OS releases. See the **protocols** statement for details on the main differences as follows:

- For PTX Series routers and QFX Series switches, see *protocols (DDoS) (PTX Series and QFX Series)*.
- For all other routing devices and EX9200 switches, see [protocols \(DDoS\)](#).

The remaining statements in this configuration statement hierarchy are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

NOTE: PTX Series routers and QFX10002-60C switches do not support the **bypass-aggregate** option.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Control Plane Distributed Denial-of-Service (DDoS) Protection Overview

Configuring Control Plane DDoS Protection

dhcp-tags (Adjustment Control Profiles)

Syntax

```
dhcp-tags {
  algorithm algorithm;
  priority priority;
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name application]
```

Release Information

Statement introduced in Junos OS Release 15.1R1.

Description

Configure the shaping rate adjustment controls for the DHCP tags application. DHCP tags are supported for session negotiation for DHCPv4 and DHCPv6 over IP demux and VLAN demux interfaces. This means that shaping rates received from DHCP option 82 in the DISCOVER message or DHCPv6 option 17 in the SOLICIT can be used to apply a shaping rate to the logical interface.

NOTE: Single-session dual-stack DHCP is not fully supported; for example, when rates vary between the individual DHCP sessions during negotiation.

Options

algorithm—Rate adjustment algorithm used by the DHCP Tags application.

Values:

- **adjust-always**—Adjust the shaping rate unconditionally.
- **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.
- **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-never**—Do not perform rate adjustments.

Default: **adjust-less**

priority—Priority of the DHCP tags application in the adjustment control profile.

Range: 1 through 10; 1 is the highest priority.

Default: 2

Required Privilege Level

interfaces—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring CoS Adjustment Control Profiles | 186](#)

[CoS Adjustment Control Profiles Overview | 183](#)

[Verifying the CoS Adjustment Control Profile Configuration | 188](#)

direction (Application Identification)

Syntax

```
direction (any | client-to-server | server-to-client);
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type signature I4-I7-signature-name  
  member member-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Specify the connection direction of the packets to which to apply pattern matching.

Options

any—Apply pattern matching to packets flowing in any direction.

client-to-server—Apply pattern matching only to packets flowing from client to server.

server-to-client—Apply pattern matching only to packets flowing from server to client.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

direction (Service Data Flow Filters)

Syntax

```
direction (uplink | downlink | both);
```

Hierarchy Level

```
[edit unified-edge pcef flow-descriptions flow-identifier],  
[edit services pcef flow-descriptions flow-identifier]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef flow-descriptions *flow-identifier*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Specify the direction in which service data flow (SDF) filters will detect service flow IP packets.

If you are using Junos OS Subscriber Aware, specify the direction at the **[edit unified-edge pcef flow-descriptions *flow-identifier*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the direction at the **[edit services pcef flow-descriptions *flow-identifier*]** hierarchy level.

Default

If you do not configure the **direction** statement, the default direction is **both**.

Options

uplink—SDF filters are applied in the uplink direction.

downlink—SDF filters are applied in the downlink direction.

both—SDF filters are applied in both the uplink and downlink directions.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Service Data Flow Filters | 412](#)[Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment](#)[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

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

[Dynamic IGMP Configuration Overview | 393](#)[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)[Disabling IGMP](#)

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

[Dynamic MLD Configuration Overview](#) | 401

Disabling MLD

download (Application Identification)

Syntax

```
download {  
    ignore-server-validation;  
}
```

Hierarchy Level

```
[edit services application-identification]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Skip server certification validation during the download of predefined Junos OS application signature packages. Validation is enabled by default.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Downloading and Installing Predefined Junos OS Application Signature Packages | 432](#)

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

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.

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 | 40](#)

[Configuring Schedulers in a Dynamic Profile for Subscriber Access | 52](#)

[scheduler \(Dynamic Scheduler Maps\) | 1134](#)

Determining Packet Drop Behavior by Configuring Drop Profile Maps for Schedulers

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 remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access](#) | 40

[Configuring Schedulers in a Dynamic Profile for Subscriber Access](#) | 52

[scheduler \(Dynamic Scheduler Maps\)](#) | 1134

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 classifiers 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 | 40](#)

[Applying a Classifier to a Subscriber Interface in a Dynamic Profile | 221](#)

[classifiers \(Definition\)](#)

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 | 40](#)

[Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile | 219](#)

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

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 classifiers 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 | 40](#)

[Applying a Classifier to a Subscriber Interface in a Dynamic Profile | 221](#)

[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 Release 10.1.

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](#) | 40

rewrite-rules

dtcp-only (System Services)

Syntax

```
dtcp-only;
```

Hierarchy Level

```
[edit system services]
```

Release Information

Statement introduced in Junos OS Release 17.3R1.

Description

Prevent RADIUS-initiated subscriber secure policy mirroring from being enabled, while allowing both DTCP-initiated mirroring and DTCP-based flow-tap services (FlowTapLite) to be enabled. This statement has no effect on existing RADIUS-initiated mirroring services. You must issue the statement before such services are activated for a subscriber. Subscriber login and session establishment are not affected.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring | 597](#)

[Subscriber Secure Policy Overview | 553](#)

[Configuring Support for Subscriber Secure Policy Mirroring | 574](#)

Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs

dynamic-class-of-service-options (Dynamic Traffic Shaping)

Syntax

```
dynamic-class-of-service-options {
    vendor-specific-tags tag;
}
```

Hierarchy Level

[edit **dynamic-profiles** *profile-name* **class-of-service**]

Release Information

Statement introduced in Junos OS Release 12.1.

apply-to-interface-set option added in Junos OS Release 19.3R1.

Description

Configure CoS attributes applied to dynamic interfaces or interface sets based on access line information received in PPPoE discovery packets on dynamic subscriber interfaces. The packets convey the information in PPPoE-IA tags.

Options

vendor-specific-tags tag—Specify one or more of the following PPPoE-IA tags to set CoS attributes:

- **access-loop-encapsulation**—Set the overhead-accounting CoS attribute based on the Access-Loop-Encapsulation tag (0x91). The tag enables the router to determine whether the access loop is cell-based or frame-based.
- **actual-data-rate-downstream**—Set the shaping-rate class-of-service attribute based on the Actual-Data-Rate-Downstream tag (0x82).
- **apply-to-interface-set**—Enable CoS to determine the node to which it conditionally applies the adjusted downstream rate. By default, when a CoS traffic control profile is configured for both a logical interface and the underlying (child) interface set, the adjusted rate values (shaping rate, overhead bytes, and overhead mode) are applied to the logical interface, according to the adjustment control profile. When you specify this option, the following conditions are evaluated:
 - If the `$junos-interface-set-name` has been resolved by AAA authorization, CoS applies the adjusted values to the underlying interface set to which the logical interface belongs. The name for this predefined variable is supplied by the Juniper Networks Qos-Set-Name VSA, 26–130. The application of the adjusted values is governed by the adjustment control profile, regardless of whether a traffic control profile is present for the logical interface.
 - If the `$junos-interface-set-name` has not been resolved by AAA authorization, CoS applies the adjusted values to the logical interface, governed by the adjustment control profile. If a traffic control profile

is not present for the logical interface, then the adjusted values are applied to the underlying interface set.

This option is appropriate for both five-level and four-level CoS hierarchies.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Setting Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags | 121](#)

[Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces | 123](#)

[Bandwidth Management for Downstream Traffic in Edge Networks Overview | 109](#)

dynamic-profiles

Syntax

```
dynamic-profiles {
  profile-name {
    class-of-service {
      dynamic-class-of-service-options {
        vendor-specific-tags tag;
      }
      interfaces {
        interface-name ;
      }
      unit logical-unit-number {
        classifiers {
          type (classifier-name | default);
        }
        output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
        report-ingress-shaping-rate bps;
        rewrite-rules {
          dscp (rewrite-name | default);
          dscp-ipv6 (rewrite-name | default);
          ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
          inet-precedence (rewrite-name | default);
        }
      }
    }
  }
  scheduler-maps {
    map-name {
      forwarding-class class-name scheduler scheduler-name;
    }
  }
  schedulers {
    (scheduler-name) {
      buffer-size (seconds | percent percentage | remainder | temporal microseconds);
      drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any | non-tcp |
        tcp) drop-profile profile-name;
      excess-priority (low | high | $junos-cos-scheduler-excess-priority);
      excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
      overhead-accounting (shaping-mode) <bytes (byte-value>;
      priority priority-level;
      shaping-rate (rate | predefined-variable);
      transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
    }
  }
}
```

```

}
traffic-control-profiles profile-name {
    adjust-minimum rate;
    delay-buffer-rate (percent percentage | rate);
    excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
    excess-rate-high (percent percentage | proportion value);
    excess-rate-low (percent percentage | proportion value);
    guaranteed-rate (percent percentage | rate) <burst-size bytes>;
    max-burst-size cells;
    overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
    peak-rate rate;
    scheduler-map map-name;
    shaping-rate (percent percentage | rate | predefined-variable) <burst-size bytes>;
    shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
    sustained-rate rate;
}
}

```

```

firewall {
  family family {
    fast-update-filter filter-name {
      interface-specific;
      match-order [match-order];
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
        only-at-create;
      }
    }
    filter filter-name {
      enhanced-mode-override;
      fast-lookup-filter;
      instance-shared;
      interface-shared;
    }
    interface-specific;
    term term-name {
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
      only-at-create;
    }
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
  hierarchical-policer uid {
    aggregate {

```

```

    if-exceeding {
        bandwidth-limit-limit bps;
        burst-size-limit bytes;
    }
    then {
        policer-action;
    }
}
premium {
    if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    then {
        policer-action;
    }
}
}
policer uid {
    filter-specific;
    if-exceeding {
        (bandwidth-limit bps | bandwidth-percent percentage);
        burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    then {
        policer-action;
    }
}
}

```

```
three-color-policer uid {  
    action {  
        loss-priority high then discard;  
    }  
    logical-interface-policer;  
    single-rate {  
        (color-aware | color-blind);  
        committed-burst-size bytes;  
        committed-information-rate bps;  
        excess-burst-size bytes;  
    }  
    two-rate {  
        (color-aware | color-blind);  
        committed-burst-size bytes;  
        committed-information-rate bps;  
        peak-burst-size bytes;  
        peak-information-rate bps;  
    }  
}  
}
```

```

interfaces interface-name {
  interface-set interface-set-name {
    interface interface-name {
      unit logical unit number {
        advisory-options {
          downstream-rate rate;
          upstream-rate rate;
        }
      }
    }
  }
}

unit logical-unit-number {
  actual-transit-statistics;
  auto-configure {
    agent-circuit-identifier {
      dynamic-profile profile-name;
    }
    line-identity {
      include {
        accept-no-ids;
        circuit-id;
        remote-id;
      }
      dynamic-profile profile-name;
    }
  }
  encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux | atm-mlppp-llc |
    atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc
    | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | frame-relay-ccc |
    frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type | frame-relay-ether-type-tcc |
    multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc |
    vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls);
  family family {
    address address;
    filter {
      adf {
        counter;
        input-precedence precedence;
        not-mandatory;
        output-precedence precedence;
        rule rule-value;
      }
      input filter-name (
        precedence precedence;

```

```

    shared-name filter-shared-name;
}
output filter-name {
    precedence precedence;
    shared-name filter-shared-name;
}
}
rpf-check {
    fail-filter filter-name;
    mode loose;
}
service {
    input {
        service-set service-set-name {
            service-filter filter-name;
        }
        post-service-filter filter-name;
    }
    input-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (push | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
    output {
        service-set service-set-name {
            service-filter filter-name;
        }
    }
    output-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (pop | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
    pcef pcef-profile-name {
        activate rule-name | activate-all;
    }
}
unnumbered-address interface-name <preferred-source-address address>;
}

```

```

filter {
    input filter-name (
        shared-name filter-shared-name;
    )
    output filter-name {
        shared-name filter-shared-name;
    }
}
host-prefix-only;
ppp-options {
    aaa-options aaa-options-name;
    authentication [ authentication-protocols ];
    chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
    }
    ignore-magic-number-mismatch;
    initiate-ncp (dual-stack-passive | ipv6 | ip)
    ipcp-suggest-dns-option;
    mru size;
    mtu (size | use-lower-layer);
    on-demand-ip-address;
    pap;
    peer-ip-address-optional;
    local-authentication {
        password password;
        username-include {
            circuit-id;
            delimiter character;
            domain-name name;
            mac-address;
            remote-id;
        }
    }
}
targeted-options {
    backup backup;
    group group;
    primary primary;
    weight ($junos-interface-target-weight | weight-value);
}

```



```

telemetry {
  subscriber-statistics;
  queue-statistics {
    interface $junos-interface-name {
      refresh rate;
      queues queue set;
    }
    interface-set $junos-interface-set-name {
      refresh rate;
      queues queue set;
    }
  }
}
vlan-id number;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
}
interfaces {
  demux0 {...}
}
interfaces {
  pp0 {...}
}
policy-options {
  prefix-list uid {
    ip-addresses;
    dynamic-db;
  }
}
predefined-variable-defaults predefined-variable <variable-option> default-value;
profile-type remote-device-service;

```

```

protocols {
  igmp {
    interface interface-name {
      accounting;
      disable;
      group-limit limit;
      group-policy;
      group-threshold value;
      immediate-leave
      log-interval seconds;
      no-accounting;
      oif-map;
      passive;
      promiscuous-mode;
      ssm-map ssm-map-name;
      ssm-map-policy ssm-map-policy-name
      static {
        group group {
          source source;
        }
      }
      version version;
    }
  }
  mld {
    interface interface-name {
      (accounting | no-accounting);
      disable;
      group-limit limit;
      group-policy;
      group-threshold value;
      immediate-leave;
      log-interval seconds;
      oif-map;
      passive;
      ssm-map ssm-map-name;
      ssm-map-policy ssm-map-policy-name;
      static {
        group multicast-group-address {
          exclude;
          group-count number;
          group-increment increment;
          source ip-address {
            source-count number;

```

```
        source-increment increment;  
    }  
}  
}  
version version;  
}  
}  
router-advertisement {  
    interface interface-name {  
        current-hop-limit number;  
        default-lifetime seconds;  
        (managed-configuration | no-managed-configuration);  
        max-advertisement-interval seconds;  
        min-advertisement-interval seconds;  
        (other-stateful-configuration | no-other-stateful-configuration);  
        prefix prefix;  
        reachable-time milliseconds;  
        retransmit-timer milliseconds;  
    }  
}  
}
```

```

routing-instances routing-instance-name {
  interface interface-name;
  routing-options {
    access {
      route prefix {
        next-hop next-hop;
        metric route-cost;
        preference route-distance;
        tag route-tag;
        tag2 route-tag2;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
    multicast {
      interface interface-name {
        no-qos-adjust;
      }
    }
  }
  rib routing-table-name {
    access {
      route prefix {
        next-hop next-hop;
        metric route-cost;
        preference route-distance;
        tag route-tag;
        tag2 route-tag2;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
  }
}

```

```

routing-options {
  access {
    route prefix {
      next-hop next-hop;
      metric route-cost;
      preference route-distance;
      tag route-tag;
      tag2 route-tag2;
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
}

multicast {
  interface interface-name {
    no-qos-adjust;
  }
}

services {
  captive-portal-content-delivery {
    auto-deactivate value;
    rule name {
      match-direction (input | input-output | output);
      term name {
        then {
          accept;
          redirect url;
          rewrite destination-address address <destination-port port-number>;
          syslog;
        }
      }
    }
  }
}

```

```

variables {
  variable-name {
    default-value default-value;
    equals expression;
    mandatory;
    uid;
    uid-reference;
  }
}

```

Hierarchy Level

[edit]

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the **filter**, **policer**, **hierarchical-policer**, **three-color-policer**, and **policy options** hierarchy levels introduced in Junos OS Release 11.4.

Description

Create dynamic profiles for use with DHCP or PPP client access.

Options

profile-name—Name of the dynamic profile; string of up to 80 alphanumeric characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring a Basic Dynamic Profile

Configuring Dynamic VLANs Based on Agent Circuit Identifier Information

Dynamic Profiles for Subscriber Management

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](#) | 124

enable-performance-mode (Application Identification)

Syntax

```
enable-performance-mode;
```

Hierarchy Level

```
[edit services application-identification]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Set the deep packet inspection (DPI) in performance mode for application identification. Performance mode improves application traffic throughput by limiting the maximum DPI processing to four packets per session. This limit includes both client-to-server and server-to-client directions. By default, performance mode is disabled.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview](#) | 431

enhanced-mode

Syntax

```
enhanced-mode;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],  
[edit firewall filter filter-name],  
[edit firewall family family-name filter filter-name],  
[edit logical-systems logical-system-name firewall filter filter-name],  
[edit logical-systems logical-system-name firewall family family-name filter filter-name]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Limit static service filters or API-client filters to term-based filter format only for inet or inet6 families when enhanced network services mode is configured at the **[edit chassis network-services]** hierarchy level. You cannot attach enhanced mode filters to local loopback, management, or MS-DPC interfaces. These interfaces are processed by the Routing Engine and DPC modules and can accept only compiled firewall filter format. In cases where both filter formats are needed for dynamic service filters, you can use the **enhanced-mode-override** statement on the specific filter definition to override the default filter term-based only format of chassis network-service enhanced IP mode. The **enhanced-mode** and the **enhanced-mode-override** statements are mutually exclusive; you can define the filter with either **enhanced-mode** or **enhanced-mode-override**, but not both.

NOTE:

For MX Series routers with MPCs, you need to initialize Trio-only match filters (that is, a filter that includes at least one match condition or action that is only supported by the Trio chipset) by walking the corresponding SNMP MIB. For example, for any filter that is configured or changed with respect to their Trio only filters, you need to run a command such as the following: **show snmp mib walk (ascii | decimal) object-id**. This forces Junos to learn the filter counters and ensure that the filter statistics are displayed. This guidance applies to all **enhanced-mode** firewall filters. It also applies to *Firewall Filter Match Conditions for IPv4 Traffic* with flexible match filter terms for offset-range or offset-mask, **gre-key**, and *Firewall Filter Match Conditions for IPv6 Traffic* with any of the following match conditions: **payload-protocol**, **extension headers**, **is_fragment**. It also applies to filters with either of the following *Firewall Filter Terminating Actions*: **encapsulate** or **decapsulate**, or either of the following *Firewall Filter Nonterminating Actions*: **policy-map**, and **clear-policy-map**.

When used with one of the chassis enhanced network services modes, firewall filters are generated in term-based format for use with MPC modules. Do not use enhanced mode for firewall filters that are intended for control plane traffic. Control plane filtering is handled by the Routing Engine kernel, which cannot use the term-based format of the enhanced mode filters.

If enhanced network services are not configured for the chassis, the **enhanced-mode** statement is ignored and any enhanced mode firewall filters are generated in both term-based and the default, compiled format. Only term-based (enhanced) firewall filters will be generated, regardless of the setting of the **enhanced-mode** statement at the **[edit chassis network-services]** hierarchy level, if any of the following are true:

- Flexible filter match conditions are configured at the **[edit firewall family family-name filter filter-name term term-name from]** or **[edit firewall filter filter-name term term-name from]** hierarchy levels.
- A tunnel header push or pop action, such as GRE encapsulate or decapsulate is configured at the **[edit firewall family family-name filter filter-name term term-name then]** hierarchy level.
- Payload-protocol match conditions are configured at the **[edit firewall family family-name filter filter-name term term-name from]** or **[edit firewall filter filter-name term term-name from]** hierarchy levels.
- An extension-header match is configured at the **[edit firewall family family-name filter filter-name term term-name from]** or **[edit firewall filter filter-name term term-name from]** hierarchy levels.
- A match condition is configured that only works with MPC cards, such as firewall bridge filters for IPv6 traffic.

For packets sourced from the Routing Engine, the Routing Engine processes Layer 3 packets by applying output filters to the packets and forwards Layer 2 packets to the Packet Forwarding Engine for transmission. By configuring the enhanced mode filter, you explicitly specify that only the term-based filter format is used, which also implies that the Routing Engine cannot use this filter.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Network Services Mode Overview</i>
Firewall Filters and Enhanced Network Services Mode Overview 359
Configuring a Filter for Use with Enhanced Network Services Mode 362
<i>Firewall Filter Match Conditions for IPv4 Traffic</i>
<i>Firewall Filter Match Conditions for IPv6 Traffic</i>
<i>Firewall Filter Terminating Actions</i>
<i>Firewall Filter Flexible Match Conditions</i>

enhanced-mode-override

Syntax

```
enhanced-mode-override;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],  
[edit firewall filter filter-name],  
[edit firewall family family-name filter filter-name],  
[edit logical-systems logical-system-name firewall filter filter-name],  
[edit logical-systems logical-system-name firewall family family-name filter filter-name]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Description

Overrides the default filter enhanced-mode of dynamic service filters when the chassis is running in network-services enhanced IP mode. It functions similarly to the **enhanced-mode** statement used to override the default IP mode of static filters when the chassis is running in network-services enhanced IP mode.

When the chassis is running in network-service enhanced IP mode, all dynamic service inet and inet6 firewall filters are automatically generated in term-based filter format only. For any dynamic service filter that must be generated in both term-based and compiled formats, you must specifically configure the **enhanced-mode-override** statement for that filter definition.

Similar to how the filter **enhanced-mode** statement functions, if the chassis is not running in network-services enhanced IP mode, then the **enhanced-mode-override** statement is ignored.

NOTE: The **enhanced-mode** and the **enhanced-mode-override** statements are mutually exclusive; you can define the filter with either **enhanced-mode** or **enhanced-mode-override**, but not both.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[enhanced-mode | 842](#)

[Network Services Mode Overview](#)

[Firewall Filters and Enhanced Network Services Mode Overview | 359](#)

[Configuring a Filter for Use with Enhanced Network Services Mode | 362](#)

enhanced-policer

Syntax

```
enhanced-policer
```

Hierarchy Level

```
[edit chassis]
```

Release Information

Statement introduced in Junos OS Release 12.3 for MX Series.

Description

Collect additional statistics to be displayed using **show** commands. An FPC restart is required after changing this configuration.

A warning log message is generated when you commit a configuration that contains the **enhanced-policer** statement. The log message states that the enhanced policer is enabled on FPCs only after they are restarted. 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](#)

[Enhanced Policer Statistics Overview | 390](#)

[show policer](#)

[show firewall | 1378](#)

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.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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

<i>Three-Color Policer Configuration Overview</i>
<i>Policer Bandwidth and Burst-Size Limits</i>
<i>Policer Color-Marking and Actions</i>
<i>Dual Token Bucket Algorithms</i>
<i>Determining Proper Burst Size for Traffic Policers</i>
committed-burst-size 788
committed-information-rate 790

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.

none option added 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 MPC Interfaces](#) | **151**
[scheduler](#) | **1134**

excess-rate (Dynamic Schedulers)

Syntax

```
excess-rate percent (percentage | $junos-cos-scheduler-excess-rate | proportion value);
```

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.

value—Proportion of the excess bandwidth to share.

Range: 0 through 1000

NOTE: The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 0 through 1000.

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](#) | 40

[Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces](#) | 151

| output-traffic-control-profile | 1006

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

NOTE: The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 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](#) | 40

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces | 151

output-traffic-control-profile | 1006

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

NOTE: The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 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](#) | 40

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces | 151

output-traffic-control-profile | 1006

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

NOTE: The proportion of excess bandwidth on MPC2-3D MPCs can be configured with increments of 1 from 0 through 1000. All other MPCs should be configured with increments of 10 from 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 | 40](#)

[Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MPC Interfaces | 151](#)

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

[Dynamic MLD Configuration Overview](#) | 401

Enabling MLD Static Group Membership

fail-filter (Dynamic Profiles)

Syntax

```
fail-filter filter-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family rpf-check],  
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family family rpf-check]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Description

Specify a filter that evaluates packets that fail a unicast RPF check. The filter determines what action to take with the failed packets. If the fail filter is not configured, the failed packets are silently discarded.

Options

filter-name—Name of the filter that evaluates packets that fail the RPF check.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring a Fail Filter for Unicast RPF in Dynamic Profiles for Subscriber Interfaces | 349](#)

[Unicast RPF in Dynamic Profiles for Subscriber Interfaces | 347](#)

Understanding Unicast RPF (Routers)

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;
    }
  }
  filter filter-name {
    enhanced-mode-override;
    fast-lookup-filter;
    instance-shared;
    interface-shared;
    interface-specific;
    term term-name {
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
    }
  }
}
```

Hierarchy Level

[edit **dynamic-profiles** *profile-name* **firewall**]

Release Information

Statement introduced in Junos OS Release 9.6.

any option added in Junos OS Release 16.1.

Description

Configure fast update filters or parameterized filters for a protocol family in a dynamic client profile or a dynamic service profile.

Options

family—Protocol family:

- **any**—Filter packets based on protocol-independent match conditions.
- **inet**—Filter Internet Protocol version 4 suite packets.
- **inet6**—filter Internet Protocol version 6 suite packets.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Fast Update Filters | 328](#)

[Parameterized Filters Overview | 249](#)

Guidelines for Configuring Firewall Filters

[Configuring Unique Identifiers for Parameterized Filters | 253](#)

family (Dynamic Standard Interface)

Syntax

```
family family {
    access-concentrator name;
    address address;
    direct-connect;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
        adf {
            counter;
            input-precedence precedence;
            not-mandatory;
            output-precedence precedence;
            rule rule-value;
        }
        input filter-name {
            precedence precedence;
            shared-name filter-shared-name;
        }
        output filter-name {
            precedence precedence;
            shared-name filter-shared-name;
        }
    }
    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> <filter [aci]>;
  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.

pppoe option added 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**—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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Static Routing on Logical Systems

Configuring the Protocol Family

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 remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

filter (Configuring)

Syntax

```
filter filter-name {
  accounting-profile name;
  enhanced-mode;
  fast-lookup-filter;
  filter-list-template;
  interface-shared;
  interface-specific;
  physical-interface-filter;
  promote gre-key;
  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 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 (" "). 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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Guidelines for Configuring Firewall Filters</i>
<i>Guidelines for Applying Standard Firewall Filters</i>
<i>Configuring Multifield Classifiers</i>
<i>Using Multifield Classifiers to Set Packet Loss Priority</i>
<i>simple-filter</i>

filter (Dynamic Profiles Filter Attachment)

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 **family any**, **family inet**, or **family inet6**. The filters can be classic filters, fast update filters, or (for the **adf** statement) Ascend-Data-Filters.

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

For general information about configuring firewall filters, see the Junos OS Routing Policies, Firewall Filters and Traffic Policers User Guide for Routing Devices .
<i>Firewall Filters Overview</i>
Understanding Dynamic Firewall Filters 226
Classic Filters Overview 229
Basic Classic Filter Syntax 232
Parameterized Filters Overview 249

filter (Dynamic Profiles Filter Creation)

Syntax

```
filter filter-name {
  enhanced-mode-override;
  fast-lookup-filter;
  instance-shared;
  interface-shared;
  interface-specific;
  term term-name {
    from {
      match-conditions;
    }
    then {
      action;
      action-modifiers;
    }
  }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family]
```

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Create firewall filters to be applied by 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 name can also be a predefined variable.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

.

filter (Dynamic Interface Unit)

Syntax

```
filter {
  input filter-name {
    shared-name filter-shared-name;
  }
  output filter-name {
    shared-name filter-shared-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 statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Guidelines for Configuring Firewall Filters

[Understanding Dynamic Firewall Filters | 226](#)

[Classic Filters Overview | 229](#)

[Basic Classic Filter Syntax | 232](#)

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

By default, a policer operates in *term-specific* mode, which means that for a given firewall filter the Junos OS creates a separate policer instance for every filter term that references the policer. You can, however, use a common policer instance for all terms within the same firewall filter by setting the *filter-specific* option in the policer. In addition, for IPv4 firewall filters with multiple terms that reference the same policer, filter-specific mode counts and monitors the activity of the policer at the firewall filter level.

NOTE: Both filter-specific and term-specific apply to prefix-specific policer sets.

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

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;
      }
    }
    filter filter-name {
      enhanced-mode-override;
      fast-lookup-filter;
      instance-shared;
      interface-shared;
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
  hierarchical-policer uid {
    aggregate {
      if-exceeding {
        bandwidth-limit-limit bps;
        burst-size-limit bytes;
      }
      then {
        policer-action;
      }
    }
  }
}

```

```

}
premium {
  if-exceeding {
    bandwidth-limit bps;
    burst-size-limit bytes;
  }
  then {
    policer-action;
  }
}
}
policer uid {
  filter-specific;
  if-exceeding {
    (bandwidth-limit bps | bandwidth-percent percentage);
    burst-size-limit bytes;
  }
  logical-bandwidth-policer;
  logical-interface-policer;
  physical-interface-policer;
  then {
    policer-action;
  }
}
three-color-policer uid {
  action {
    loss-priority high then discard;
  }
  logical-interface-policer;
  single-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    peak-burst-size bytes;
    peak-information-rate bps;
  }
}
}
}

```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Configure firewall filters and policers in a dynamic client profile or a dynamic service profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Methods for Regulating Traffic by Applying Hierarchical Policers | 369](#)

[Configuring Fast Update Filters | 328](#)

flow-descriptions

Syntax

```
flow-descriptions flow-identifier {
  direction (uplink | downlink | both);
  local-port-range {
    low lower-boundary high upper-boundary;
  }
  local-ports number;
  no-send-to-ue;
  protocol protocol-number;
  remote-address (ipv4-address ipv4-address | ipv6-address ipv6-address);
  remote-port-range {
    low lower-boundary high upper-boundary;
  }
  remote-ports number;
}
```

Hierarchy Level

```
[edit unified-edge pcef],
[edit services pcef]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Specify a service data flow (SDF) filter (flow identifier) that includes one or more filtering parameters (address, protocol, and port) to identify the subscriber traffic that you want the SDF filter to detect. SDF filters are specified in a PCC rule to identify the Layer 3 or Layer 4 IP packet flows that you want to receive a particular treatment.

NOTE: A PCC rule must include at least one SDF filter and can include a maximum of 15 SDF filters.

If you are using Junos OS Subscriber Aware, specify the name of the SDF filter at the **[edit unified-edge pcef]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the SDF filter at the **[edit services pcef]** hierarchy level.

Options

flow-identifier—Name of the SDF filter.

Range: 1 through 63 characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Service Data Flow Filters | 412](#)

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

flow-tap

Syntax

```
flow-tap {
  (interface interface-name | tunnel-interface interface-name);
  family (inet | inet6);
}
```

Hierarchy Level

```
[edit services]
```

Release Information

Statement introduced in Junos OS Release 8.1.

ccc option introduced in Junos OS Release 17.2.

Description

Enable the flow-tap service or FlowTapLite service on an interface. FlowTapLite is a lighter version of the flow-tap application that is available only on tunnel interfaces on MX Series platforms, M120 Series routers, and M320 Series routers with Enhanced III FPCs only.

Starting in Junos OS Release 17.3R1, the FlowTapLite service can run concurrently with the radius-flow-tap service on the same MX Series router. The radius-flow-tap service (**[edit services radius-flow-tap]**) is required for subscriber secure policy mirroring on MX Series routers.

In earlier releases, the FlowTapLite and radius-flow-tap services cannot run concurrently on an MX Series router, which prevents you from running FlowTapLite monitoring and subscriber secure policy mirroring at the same time.

Options

interface *interface-name*—Use the specified interface for the flow-tap application.

tunnel-interface *interface-name*—Use the specified tunnel interface for the FlowTapLite application.

family—(Not applicable for FlowTapLite) Apply flow-tap services to the specified family. If you do not specify an option, the flow-tap service is applied only to IPv4 traffic.

- **inet**—IPv4 traffic.
- **inet6**—IPv6 traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release History Table

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, the FlowTapLite service can run concurrently with the radius-flow-tap service on the same MX Series router.

RELATED DOCUMENTATION

| *Configuring Junos Packet Vision on MX, M and T Series Routers*

flow-tap-dtcp

Syntax

```
flow-tap-dtcp {
  ssh {
    connection-limit limit;
    rate-limit limit;
  }
}
```

Hierarchy Level

[edit system services]

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Configure Dynamic Tasking Control Protocol (DTCP) sessions to run over SSH in support of the flow-tap, FlowTapLite, or radius-flow-tap services. Note that the flow-tap feature is not supported on outbound, or egress, traffic. Only inbound, or ingress, traffic is supported.

This statement is required for DTCP-initiated subscriber secure policy mirroring (radius-flow-tap service).

Options

connection-limit *limit*—(Optional) Maximum number of connections allowed.

Range: 1 through 250

Default: 75

rate-limit *limit*— (Optional) Maximum number of connection attempts allowed per minute.

Range: 1 through 250

Default: 150

Required Privilege Level

flow-tap— To view this statement in the configuration.

flow-tap-control— To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring a DTCP-over-SSH Connection to the Mediation Device | 596](#)

Configuring Flow-Tap Security Properties on MX, M and T Series Routers

Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs

Configuring DTCP-over-SSH Service for the Flow-Tap Application

flows (PCC Rules)

Syntax

```
flows ([flow-identifier] | any);
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rule-name from],  
[edit services pcef pcc-rules rule-name from]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-rules *rule-name* from]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the service data flow (SDF) filters (flow identifiers) that define the match criteria for the policy and charging control (PCC) rule. You can configure a maximum of 15 SDF filters. You must include the **flows** statement in a PCC rule. If you do not want to filter subscriber traffic based on SDF filters, use the **any** option.

If you are using Junos OS Subscriber Aware, specify the name of the SDF filter at the **[edit unified-edge pcef pcc-rules *rule-name* from]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the SDF filter at the **[edit services pcef pcc-rules *rule-name* from]** hierarchy level.

Options

flow-identifier—Name of an SDF filter that is used to detect IP packet flows. You can configure a maximum of 15 SDF filters. The referenced SDF filters must be configured.

Range: 1 through 63 characters.

any—All IP packet flows.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Policy and Charging Control Rules | 418](#)

[Configuring Service Data Flow Filters | 412](#)

format (LRF Profile)

Syntax

```
format ipfix;
```

Hierarchy Level

```
[edit services lrf profile profile-name template template-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure a format for the template. Only the IPFIX format is supported for this release.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management | 459](#)

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](#) | 40

[Configuring Schedulers in a Dynamic Profile for Subscriber Access](#) | 52

forwarding-class (PCC Action Profiles)

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],  
[edit services pcef pcc-action-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the forwarding class to which packets must be assigned.

If you are using Junos OS Subscriber Aware, specify the forwarding class at the **[edit unified-edge pcef pcc-action-profiles *profile-name*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the forwarding class at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level.

Options

class-name—Name of the forwarding class.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

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 | 553](#)[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

fpc (MX Series 5G Universal Routing Platforms)

Syntax

```
fpc slot-number {
  inline-services {
    flow-table-size {
      ipv4-flow-table-size units;
      ipv4-flow-table-size units;
      ipv6-extended-attrib;
    }
  }
  ir-mode (R | IR);
  pic number {
    inline-services {
      bandwidth (1g | 10g);
    }
    port-mirror-instance port-mirroring-instance-name-pic-level;
    tunnel-services {
      bandwidth (1g | 10g)
    }
  }
  port-mirror-instance port-mirroring-instance-name-fpc-level;
}
```

Hierarchy Level

[edit chassis]

Release Information

Statement introduced in Junos OS Release 8.2.

port-mirror-instance option added in Junos OS Release 9.3.

ipv6-extended-attrib option added in Junos OS Release 14.2 for MX Series routers.

Description

Configure properties for the DPC or MPC and corresponding Packet Forwarding Engines to create tunnel interfaces.

(MX Series Virtual Chassis only) When you configure chassis properties for MPCs installed in a Virtual Chassis member router, statements included at the **[edit chassis member member-id fpc slot slot-number]** hierarchy level apply to the MPC in the specified slot number only on the specified member router in the Virtual Chassis. Statements included at the **[edit chassis fpc slot slot-number]** hierarchy level apply to the MPCs in the specified slot number on *each* member router in the Virtual Chassis.

BEST PRACTICE: To ensure that the statement you use to configure MPC chassis properties in an MX Series Virtual Chassis applies to the intended member router and MPC, we recommend that you always include the **member member-ID** option before the **fpc** statement, where **member-id** is 0 or 1 for a two-member MX Series Virtual Chassis.

Options

fpc slot-number—Specify the slot number of the DPC.

Range: 0 through 11

pic number—Specify the number of the Packet Forwarding Engine. Each DPC includes four Packet Forwarding Engines.

Range: 0 through 4

port-mirror-instance port-mirroring-instance-name-fpc-level—Associate a port-mirroring instance with the DPC and its corresponding PICs. The port-mirroring instance is configured under the **[edit forwarding-options port-mirroring]** hierarchy level.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Port-Mirroring Instances on MX Series 5G Universal Routing Platforms

Enabling Inline Service Interfaces

Virtual Chassis Components Overview

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 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-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 183
Configuring CoS Adjustment Control Profiles 186
adjustment-control-profiles 712
Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates 111
Bandwidth Management for Downstream Traffic in Edge Networks Overview 109
<i>egress-shaping-overhead</i>
bytes 756
cell-mode 763

from (Captive Portal Content Delivery Tags)

Syntax

```
from {
  destination-address address <except>;
}
```

Syntax

Hierarchy Level

```
[edit services captive-portal-content-delivery rule rule-name term term-name]
```

Release Information

Statement introduced in Junos OS Release 19.1R1 on MX Series routers.

Description

Specify one or more addresses in a CPCD rule term to identify traffic that can have header tags inserted into GET messages. Tag insertion is supported only for Routing-Engine-based, static CPCD.

Options

address—Destination IPv4 or IPv6 address for traffic to be tagged.

except—(Optional) Exclude the specified address 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

[Configuring Routing Engine-Based, Static HTTP Redirect Services](#) | 505

from (PCC Rules)

Syntax

```
from {
  <application-groups [application-group-name]>;
  <applications [application-name]>;
  flows ([flow-identifier] | any);
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rule-name],
[edit services pcef pcc-rules rule-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-rules rule-name]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the match criteria for the policy and charging control (PCC) rules. Any referenced SDF filter, application, or application group in the **from** statement must be configured.

If you are using Junos OS Subscriber Aware, specify the match criteria at the **[edit unified-edge pcef pcc-rules rule-name]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the match criteria at the **[edit services pcef pcc-rules rule-name]** hierarchy level.

NOTE: You must include the **flows** statement. If you do not want to filter subscriber traffic based on service data flow (SDF) filters, use **flows any**.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Policy and Charging Control Rules](#) | 418

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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](#) | 553

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview](#) | 572

gate-status

Syntax

```
gate-status (uplink | downlink | uplink-downlink | disable-both);
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],  
[edit services pcef pcc-action-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Configure the gate status in a PCC action profile to enable or disable the forwarding of service flow packets. The gate status determines whether the uplink and downlink gates are opened or closed.

If you are using Junos OS Subscriber Aware, configure the gate status at the **[edit unified-edge pcef pcc-action-profiles *profile-name*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the gate status at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level.

Default

By default, if this statement is not configured, forwarding of service data flow packets is enabled in both the uplink and downlink directions.

Options

disable-both—Disable forwarding of service data flow packets in the uplink and downlink directions.

downlink—Enable forwarding of service data flow packets in the downlink direction.

uplink-downlink—Enable forwarding of service data flow packets in the uplink and downlink directions.

uplink—Enable forwarding of service data flow packets in the uplink direction.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Enabling IGMP Static Group Membership

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 **mlld 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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Dynamic MLD Configuration Overview](#) | 401

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

[Dynamic MLD Configuration Overview](#) | 401

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

[Dynamic MLD Configuration Overview](#) | 401

[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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

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

[Dynamic MLD Configuration Overview | 401](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Configuring the Number of MLD Multicast Group Joins on Logical Interfaces

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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Filtering Unwanted IGMP Reports at the IGMP Interface Level

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

[Dynamic MLD Configuration Overview](#) | 401

Filtering Unwanted MLD Reports at the MLD Interface Level

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 and optionally a burst size for the guaranteed rate.

NOTE: The guaranteed-rate burst size must not exceed the shaping-rate burst size.

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 6,400,000,000,000 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 | 40](#)

[Configuring Traffic Scheduling and Shaping for Subscriber Access | 47](#)

[output-traffic-control-profile | 1006](#)

hierarchical-policer

List of Syntax

[Syntax \(M Series, MX Series, T Series - Bandwidth-Based\) on page 911](#)

[Syntax \(MX Series - Packets-Per-Second \(pps\)-Based\) on page 911](#)

Syntax (M Series, MX Series, T Series - Bandwidth-Based)

```
hierarchical-policer hierarchical-policer-name | uid {
  aggregate {
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      discard;
    }
  }
  premium {
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      discard;
    }
  }
}
```

Syntax (MX Series - Packets-Per-Second (pps)-Based)

```
hierarchical-policer hierarchical-policer-name | uid {
  aggregate {
    if-exceeding-pps {
      pps-limit pps;
      packet-burst packets;
    }
    then {
      discard;
    }
  }
  premium {
    if-exceeding-pps (Hierarchical Policer) {
      pps-limit (Hierarchical Policer) pps;
    }
  }
}
```

```

        packet-burst (Hierarchical Policer) packets;
    }
    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 profile-name firewall]` hierarchy level introduced in Junos OS Release 11.4.

Support for **if-exceeding-pps** statement on MX Series routers with MPCs introduced in Junos OS Release 15.2.

Description

Use a hierarchical policer to rate-limit ingress Layer 2 traffic at a physical or logical interface and apply different policing actions based on whether the packets are classified as **premium** for expedited forwarding (EF) or **aggregate** for a lower priority. The two policers defined within the hierarchical policer are **aggregate** and **premium**.

Hierarchical policers are supported on Enhanced Intelligent Queuing (IQE) PICs and SONET interfaces hosted on the M120 and M320 with incoming Flexible PIC Concentrators (FPCs) as SFPC and outgoing FPCs as FFPC; on MPCs hosted on MX Series routers; on the T320, T640, and T1600 with Enhanced Intelligent Queuing (IQE) PICs; and on the T4000 with Type 5 FPC and Enhanced Scaling Type 4 FPC.

NOTE:

- The **if-exceeding-pps** statement is only supported on MX Series routers with MPCs.
- The **if-exceeding** and **if-exceeding-pps** statements are mutually exclusive and, therefore, cannot be applied at the same time.

You can configure the policer in static firewall filters or dynamic firewall filters in a dynamic client profile or a dynamic service profile.

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 the name in quotation marks (" ").

uid—When you configure a hierarchical policer at the **[edit dynamic-profiles profile name firewall]** hierarchy level, you must assign a variable UID as the policer name.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Hierarchical Policers Configuration Overview</i>
<i>Hierarchical Policers</i>
aggregate (Hierarchical Policers) 718
<i>bandwidth-limit (Hierarchical Policers)</i>
burst-size-limit (Hierarchical Policers) 752
<i>pps-limit (Hierarchical Policers)</i>
<i>packet-burst (Hierarchical Policers)</i>
if-exceeding (Hierarchical Policers) 920
<i>if-exceeding-pps (Hierarchical Policers)</i>
premium (Hierarchical Policers) 1055

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)

Syntax

```
hierarchical-scheduler {
    implicit-hierarchy;
    maximum-hierarchy-levels number;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Release Information

Statement introduced in Junos OS Release 10.1.

implicit-hierarchy option added in Junos OS Release 13.1.

Support on GRE tunnel interfaces configured on physical interfaces on MICs or MPCs in MX Series routers added in Junos OS Release 13.3.

Support for up to four hierarchy levels added in Junos OS Release 16.1.

Description

Configure hierarchical scheduling options on the interface.

The statement is supported on the following interfaces:

- MIC and MPC interfaces in MX Series routers
- GRE tunnel interfaces configured on physical interfaces hosted on MIC or MPC line cards in MX Series routers

To enable hierarchical scheduling on MX Series routers, configure the **hierarchical-scheduler** statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.

Options

implicit-hierarchy—Configure four-level hierarchical scheduling. When you include the **implicit-hierarchy** option, a hierarchical relationship is formed between the CoS scheduler nodes at level 1, level 2, level 3, and level 4. The **implicit-hierarchy** option is supported only on MPC/MIC subscriber interfaces and interface sets on MX Series routers.

maximum-hierarchy-levels *number*—Specify the maximum number of hierarchical scheduling levels allowed for node scaling, from 2 through 4 levels. The default number of levels is 3. The **maximum-hierarchy-levels** option is supported on MPC/MIC or EQ DPC subscriber interfaces and interface sets on MX Series routers.

- If you set **maximum-hierarchy-levels** to **2**, interface sets are not allowed. In this case, if you configure a level 2 interface set, you generate Packet Forwarding Engine errors.
- If you do not include the **maximum-hierarchy-levels** option, keeping the default number of hierarchy levels at 3, interface sets can be at either level 2 or level 3, depending on whether the member logical interfaces within the interface set have a traffic control profile. If any member logical interface has a traffic control profile, then the interface set is a level 2 CoS scheduler node. If no member logical interface has a traffic control profile, the interface set is at level 3.



CAUTION: MPC3E, 32x10GE MPC4E, and 2x100GE + 8x10GE MPC4E MPCs support only two levels of scheduling hierarchy. When enabling hierarchical scheduling on these cards, you must explicitly set **maximum-hierarchy-levels** to **2**.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Hierarchical CoS for Subscriber Interfaces

Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links

Configuring Hierarchical Schedulers for CoS

Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface

Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview

http-log-multiple-transactions (LRF Profile)

Syntax

```
http-log-multiple-transactions;
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure HTTP transaction logging to generate and send HTTP metadata for each transaction of a data session. This option is only relevant if the template specified in an LRF rule includes **http** in the **template-type**.

By default, HTTP transaction logging is disabled, and the HTTP transaction records for a TCP session are sent together as one group of records.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management | 459](#)

icmp-mapping (Application Identification)

Syntax

```
icmp-mapping {  
  code icmp-code;  
  order order;  
  order-priority (high | low);  
  type icmp-type;  
}
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Match Internet Control Message Protocol (ICMP) messages identified by unique code and type. This classification is intended to identify and differentiate various types of ICMP messages.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

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 classifiers ieee-802.1] hierarchy level.

default—The default mapping.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[Applying a Classifier to a Subscriber Interface in a Dynamic Profile | 221](#)

[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 statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile | 219](#)

rewrite-rules

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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 \(Hierarchical Policer\) | 718](#)

bandwidth-limit (Hierarchical Policer)

[burst-size-limit \(Hierarchical Policer\) | 752](#)

hierarchical-policer | 911

premium (Hierarchical Policer) | 1055

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure rate limits for a single-rate two-color policer.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Two-Color Policer Configuration Overview</i>
<i>Hierarchical Policer Configuration Overview</i>
<i>Basic Single-Rate Two-Color Policers</i>
<i>Bandwidth Policers</i>
<i>Filter-Specific Counters and Policers</i>
<i>Prefix-Specific Counting and Policing Actions</i>

Multifield Classification

Policer Overhead to Account for Rate Shaping in the Traffic Manager

Hierarchical Policers

igmp (Dynamic Profiles)

Syntax

```
igmp {
  interface interface-name {
    accounting;
    disable;
    group-limit policy-name;
    group-policy;
    group-threshold value;
    immediate-leave;
    log-interval seconds;
    no-accounting;
    oif-map map-name;
    passive <allow-receive> <send-general-query> <send-group-query>;
    promiscuous-mode;
    ssm-map ssm-map-name;
    ssm-map-policy ssm-map-policy-name
    static {
      group group {
        source source;
      }
    }
    version version;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* 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.

You can configure IGMP in a dynamic client profile or a dynamic service profile.

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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Understanding IGMP

Enabling IGMP

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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Specifying Immediate-Leave Host Removal for IGMP

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

[Dynamic MLD Configuration Overview](#) | 401

Specifying Immediate-Leave Host Removal for MLD

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

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 classifiers inet-precedence]** 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](#) | 40

[Applying a Classifier to a Subscriber Interface in a Dynamic Profile](#) | 221

[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 | 40](#)

[Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile | 219](#)

[rewrite-rules](#)

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

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.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Dynamic Service Sets Overview](#) | 365

[Associating Service Sets with Interfaces in a Dynamic Profile](#) | 366

inputs (Analytics)

Syntax

```
inputs {
  analytics {
    parameters {
      generate-tags value;
      sample-frequency value;
      sensors path;
    }
  }
  input-ipfix {
    parameters {
      maximum-connections number;
      tcp-port port-number;
      vrf-name name;
    }
  }
  input-jti-ipfix {
    parameters {
      record-group group-name {
        record ipfix-record-name;
        reporting-interval seconds;
      }
    }
  }
}
```

Hierarchy Level

```
[edit services analytics agent service-agents agent-name]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.


Statement introduced in Junos OS Release 18.4R1 on PTX Series routers.

input-jti-ipfix option added in Junos OS Release 18.4R1 on MX Series routers.

analytics option added in Junos OS Release 18.4R1 on MX Series and PTX Series routers.

Description

Configure parameters for a Network Telemetry Framework (NTF) service agent input plug-in. For each service agent instance, you can configure more than one input plug-in to push data to the output plug-in.



NOTE: When you modify the input plug-in configuration of a service agent instance, the associated service agent daemon is restarted.

Options

analytics parameters—Configure parameters to collect data from Junos Telemetry Interface (JTI) sensors.

generate-tags value—(Optional) Enable tag generation.

Default: Enabled

sample-frequency value—Specify the frequency interval (in seconds) at which the JTI sensor generates data to export to the data collector. Range is from 0 to 24 hours.

Default: 5 seconds

sensors file-path—Specify the resource string associated with the JTI sensor for collecting JTI data from a specific resource. The format is a file path and must be entered exactly. For a list of available JTI resource string options, see the *sensor* configuration statement and *Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)* documentation.

input-ipfix parameters—Configure parameters for the IPFIX mediation service agent to gather and consolidate IPFIX records from downstream devices.

NOTE: Any change you make to an existing **input-ipfix** plug-in configuration restarts the IPFIX service agent daemon to apply the changes.

NOTE: Although each of the parameters has a default value, you must configure at least one of the parameters to enable the plug-in. If you configure only one parameter and want to use the default value, you must specify that value.

maximum-connections number—(Optional) Maximum number of TCP connections that the IPFIX mediator can support.

Range: 1 through 500

Default: 100

tcp-port port-number—(Optional) TCP port on the IPFIX mediator that receives TCP packets; the listening port.

Default: 4739

vrf-name name—(Optional) Name of the VRF (routing instance) in which IPFIX packets are accepted.

Default: default

input-jti-ipfix parameters—Configure parameters for the IPFIX mediation service agent to collect and report local sensor data from the BNG configured as an IPFIX mediator. For each group of records, the plug-in subscribes to the specific sensor data sets associated with each record.

When you remove a record group from the configuration, the sensor sets for the member records are unsubscribed. The template IDs for the associated IPFIX records are returned to the pool for re-use.

record *ipfix-record-name*—One of the following individual IPFIX records associated with a nonconfigurable set of local sensor data. See [“Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector” on page 677](#) for the sensors collected by each record.

address-pool-utilization	port-statistics
chassis-inventory	resource-utilization
chassis-power	subscriber-statistics
dhcpv4-server-stats	thermal
interface-metadata	uptime

record-group *group-name*—Name of a group of IPFIX records that subscribes to the sensor data sets associated with the individual records that comprise the record group. You can configure a maximum of 10 record groups.

reporting-interval *seconds*—(Optional) Interval in seconds between reports for the subscribed sensor data. The interval applies to all records (and all sensor sets) in the record group.

Range: 60 through 86,400 seconds

Default: 900 seconds

Required Privilege Level

system

RELATED DOCUMENTATION

[Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data | 672](#)

[Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator | 680](#)

[Configuring NTF Agent](#)

[IPFIX Mediation on the BNG | 667](#)

[Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector | 677](#)

interface (Dynamic IGMP)

Syntax

```
interface interface-name {
  accounting;
  disable;
  distributed;
  group-limit limit;
  group-policy;
  group-threshold value;
  immediate-leave
  log-interval seconds;
  no-accounting;
  oif-map;
  passive;
  promiscuous-mode;
  ssm-map ssm-map-name;
  ssm-map-policy ssm-map-policy-name
  static {
    group group {
      source source;
    }
  }
  version version;
}
```

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Enabling IGMP

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 statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Defining ACI Interface Sets

[Guidelines for Configuring Dynamic CoS for Subscriber Access](#) | 40

[Configuring an Interface Set of Subscribers in a Dynamic Profile](#) | 192

Agent Circuit Identifier-Based Dynamic VLANs Overview

interface (Dynamic MLD)

Syntax

```

interface interface-name {
  (accounting | no-accounting);
  disable;
  distributed;
  group-limit limit;
  group-policy;
  group-threshold value;
  immediate-leave;
  log-interval seconds;
  oif-map;
  passive;
  ssm-map ssm-map-name;
  ssm-map-policy ssm-map-policy-name;
  static {
    group multicast-group-address {
      exclude;
      group-count number;
      group-increment increment;
      source ip-address {
        source-count number;
        source-increment increment;
      }
    }
  }
  version version;
}

```

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Dynamic MLD Configuration Overview](#) | 401

Enabling MLD

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 *Types of Interfaces Overview*.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Dynamic Access Routes for Subscriber Management](#)

[Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers](#)

interface-service (Services Interfaces)

Syntax

```
interface-service {  
  load-balancing-options {  
    hash-keys {  
      egress-key (destination-ip | source-ip);  
      ingress-key (destination-ip | source-ip);  
    }  
  }  
  service-interface name;  
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the device name for the interface service Physical Interface Card (PIC).

Options

service-interface *name*—Name of the service device associated with the interface-wide service set.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring Service Sets to be Applied to Services Interfaces*

interface-set (Dynamic Profiles)

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.

\$junos-phy-ifd-interface-set-name option added in Junos OS Release 16.1.

\$junos-pon-id-interface-set-name option added in Junos OS Release 17.2R1.

Description

For MX Series routers with enhanced queuing DPCs or MPC/MIC modules, configure an interface set for dynamic CoS.

Options

interface-set-name—Name of the interface set 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-phy-ifd-interface-set-name—Locally generated interface set name associated with the underlying physical interface in a dynamic profile. This predefined variable enables you to group all the subscribers on a specific physical interface so that you can apply services to the entire group of subscribers.

Another use case for this predefined variable is to conserve CoS resources in a mixed business and residential topology by collecting the residential subscribers into an interface set associated with the physical interface, so that a level 2 node is used for the interface set rather than for each residential interface. Otherwise, because the business and residential subscribers share the same interface and business subscribers require three levels of CoS, then three levels are configured for each residential subscriber. That results in an unnecessary level 2 node being consumed for each residential connection, wasting CoS resources.

- \$junos-pon-id-interface-set-name—Locally generated interface set name extracted from the DHCPv4 (Option 82, suboption 2) or DHCPv6 (Option 37) agent remote ID string inserted by an optical line

terminal (OLT) in a passive optical network (PON). The OLT must format the agent remote ID string with a pipe symbol (|) as the delimiter between substrings. The substring extracted for the interface set name consists of the characters following the last delimiter in the agent remote ID string.

The extracted substring identifies individual customer circuits in the PON to be aggregated into the interface set. You determine the format and contents of the substring, and configure your OLT to insert the information. Typically, the substring may include the name and port of the OLT accessed by the CPE optical network terminal (ONT).

- `$junos-svlan-interface-set-name`—Locally generated interface set name for use by dual-tagged VLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is ***physical_interface_name - outer_VLAN_tag***.
- `$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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[CoS for Interface Sets of Subscribers Overview | 189](#)

[Configuring an Interface Set of Subscribers in a Dynamic Profile | 192](#)

[Example: Configuring a Dynamic Service VLAN Interface Set of Subscribers in a Dynamic Profile | 211](#)

[Extracting an Option 82 or Option 37 Substring to Create an Interface Set](#)

interface-shared

Syntax

```
interface-shared;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],  
[edit 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.

NOTE: A firewall filter cannot be both interface-specific and interface-shared.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Interface-Shared Filters Overview | 289](#)

[Understanding Dynamic Firewall Filters | 226](#)

[Classic Filters Overview | 229](#)

[Basic Classic Filter Syntax | 232](#)

interface-specific (Dynamic Firewalls)

Syntax

```
interface-specific;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family 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.

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 | 328](#)

Interface-Specific Firewall Filter Instances Overview

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);
      report-ingress-shaping-rate;
      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 in a dynamic client profile or a dynamic service profile.

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile | 217](#)

interfaces (Static and Dynamic Subscribers)

Syntax

```

interfaces {
  interface-name {
    unit logical-unit-number {
      actual-transit-statistics;
      auto-configure {
        agent-circuit-identifier {
          dynamic-profile profile-name;
        }
        line-identity {
          include {
            accept-no-ids;
            circuit-id;
            remote-id;
          }
          dynamic-profile profile-name;
        }
      }
    }
    family family {
      access-concentrator name;
      address address;
      direct-connect;
      duplicate-protection;
      dynamic-profile profile-name;
      filter {
        adf {
          counter;
          input-precedence precedence;
          not-mandatory;
          output-precedence precedence;
          rule rule-value;
        }
        input filter-name {
          precedence precedence;
          shared-name filter-shared-name;
        }
        output filter-name {
          precedence precedence;
          shared-name filter-shared-name;
        }
      }
      max-sessions number;
    }
  }
}

```

```

max-sessions-vsa-ignore;
rpf-check {
    mode loose;
}
service {
    input {
        service-set service-set-name {
            service-filter filter-name;
        }
        post-service-filter filter-name;
    }
    output {
        service-set service-set-name {
            service-filter filter-name;
        }
    }
}
service-name-table table-name
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>;
unnumbered-address interface-name <preferred-source-address address>;
}
filter {
    input filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
    output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
}
host-prefix-only;
ppp-options {
    chap;
    pap;
}
proxy-arp;
service {
    pcef pcef-profile-name {
        activate rule-name | activate-all;
    }
}
}

```



```

    targeted-options {
        backup backup;
        group group;
        primary primary;
        weight ($junos-interface-target-weight | weight-value);
    }
    vlan-id;
    vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
vlan-tagging;
}
interface-set interface-set-name {
    interface interface-name {
        unit logical unit number {
            advisory-options {
                downstream-rate rate;
                upstream-rate rate;
            }
        }
    }
}
pppoe-underlying-options {
    max-sessions number;
}
}

```

```

demux0 {
  unit logical-unit-number {
    demux-options {
      underlying-interface interface-name
    }
    family family {
      access-concentrator name;
      address address;
      direct-connect;
      duplicate-protection;
      dynamic-profile profile-name;
      demux-source {
        source-prefix;
      }
      filter {
        input filter-name (
          precedence precedence;
          shared-name filter-shared-name;
        )
        output filter-name {
          precedence precedence;
          shared-name filter-shared-name;
        }
      }
      mac-validate (loose | strict);
      max-sessions number;
      max-sessions-vsa-ignore;
      rpf-check {
        fail-filter filter-name;
        mode loose;
      }
      service-name-table table-name
      short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>;
      unnumbered-address interface-name <preferred-source-address address>;
    }
    filter {
      input filter-name;
      output filter-name;
    }
    vlan-id number;
    vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
  }
}

```

```

pp0 {
  unit logical-unit-number {
    keepalives interval seconds;
    no-keepalives;
    pppoe-options {
      underlying-interface interface-name;
      server;
    }
    ppp-options {
      aaa-options aaa-options-name;
      authentication [ authentication-protocols ];
      chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
      }
      ignore-magic-number-mismatch;
      initiate-ncp (dual-stack-passive | ipv6 | ip)
      ipcp-suggest-dns-option;
      mru size;
      mtu (size | use-lower-layer);
      on-demand-ip-address;
      pap;
      peer-ip-address-optional;
      local-authentication {
        password password;
        username-include {
          circuit-id;
          delimiter character;
          domain-name name;
          mac-address;
          remote-id;
        }
      }
    }
  }
  family inet {
    unnumbered-address interface-name;
    address address;
    service {
      input {
        service-set service-set-name {
          service-filter filter-name;
        }
        post-service-filter filter-name;
      }
    }
  }
}

```

```

    output {
        service-set service-set-name {
            service-filter filter-name;
        }
    }
}
filter {
    input filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
    output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
}
}
}
}
}
stacked-interface-set {
    interface-set-name interface-set-name {
        interface-set-name interface-set-name;
    }
}
}
}

```

Hierarchy Level

[edit **dynamic-profiles** *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Define interfaces for dynamic client profiles.

Options

interface-name—The interface variable (**\$junos-interface-ifd-name**). The interface variable is dynamically replaced with the interface the DHCP client accesses when connecting to the router.

NOTE: Though we do not recommend it, you can also enter the specific name of the interface you want to assign to the dynamic profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

Configuring Dynamic PPPoE Subscriber Interfaces

Configuring Dynamic VLANs Based on Agent Circuit Identifier Information

DHCP Subscriber Interface Overview

Subscribers over Static Interfaces Configuration Overview

Demultiplexing Interface Overview

ip-protocol-mapping (Application Identification)

Syntax

```
ip-protocol-mapping {
  order order;
  order-priority (high | low);
  protocol (http | ssl | tcp | udp)
}
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

For IP traffic, identify an application by matching the IP protocol. This parameter is used to identify an application based on IP and is intended only for IP traffic.

Options

protocol-number—Industry-standard numeric protocol value.

Range: 0 through 254

You can find a complete list of industry standard protocol numbers at [the IANA website](#).

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

ipv4-address (Steering Path)

Syntax

```
ipv4-address ipv4-address;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name steering path]
[edit services pcef pcc-action-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the IPv4 address of a third-party server to which the PCC action profile steers HTTP traffic for applying services. If you configure this, the PCC action profile can only be used in PCC rules that match only HTTP-based applications and all flows.

Options

ipv4-address *ipv4-address* —Use the specified IPv4 address of the server.

Required Privilege Level

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

[Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment](#)

[Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware](#)

ipv6-address (Steering Path)

Syntax

```
ipv6-address ipv6-address;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name steering path]
[edit services pcef pcc-action-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the IPv6 address of a third-party server to which the PCC action profile steers HTTP traffic for applying services. If you configure this, the PCC action profile can only be used in PCC rules that match only HTTP-based applications and all flows.

Options

ipv6-address *ipv6-address* —Use the specified IPv6 address of the server.

Required Privilege Level

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

[Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment](#)

[Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware](#)

keep-existing-steering

Syntax

```
keep-existing-steering;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name steering],  
[edit services pcef pcc-action-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify that the PCC action profile steering attributes that a PCC rule applies at the start of a data flow will continue to be applied to that data flow when the PCC rule match conditions are modified, deleted, or added to.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

[Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment](#)

[Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware](#)

local-port-range

Syntax

```
local-port-range {  
    low low-value;  
    high high-value;  
}
```

Hierarchy Level

```
[edit unified-edge pcef flow-descriptions flow-identifier],  
[edit services pcef flow-descriptions flow-identifier]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef flow-descriptions *flow-identifier*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Specify the port range to identify the subscriber traffic that you want the service data flow (SDF) filter to detect.

NOTE: You can specify either **local-port-range** or a list of ports with **local-ports**, but not both.

If you are using Junos OS Subscriber Aware, specify the port range at the **[edit unified-edge pcef flow-descriptions *flow-identifier*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the port range at the **[edit services pcef flow-descriptions *flow-identifier*]** hierarchy level.

Default

If the **local-port-range** statement is not configured, the default is any range of local ports.

Options

low-value— Lower boundary for the port range.

Range: 1 through 65,535

high-value — Upper boundary for the port range.

Range: 1 through 65,535

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Service Data Flow Filters | 412](#)

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

local-ports

Syntax

```
local-ports [number];
```

Hierarchy Level

```
[edit unified-edge pcef flow-description flow-identifier],  
[edit services pcef flow-description flow-identifier]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef flow-description *flow-identifier*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Specify a port number or list of port numbers to identify the subscriber traffic that you want the service data flow (SDF) filter to detect.

NOTE: You can specify either a list of ports or a port range, but not both.

If you are using Junos OS Subscriber Aware, specify the port numbers at the **[edit unified-edge pcef flow-description *flow-identifier*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the port numbers at the **[edit services pcef flow-description *flow-identifier*]** hierarchy level.

Default

If the **local-ports** statement is not configured, the default is any local ports.

Options

number—Number of a port or list of port numbers. You can specify a maximum of three port numbers (separated by a space) in a list.

Range: 1 through 65,535

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Service Data Flow Filters | 412](#)

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

logging-rule (PCC Action Profile)

Syntax

```
logging-rule lrf-rule-name;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],  
[edit services pcef pcc-action-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Assign the LRF rule to the PCC action profile of a static PCC rule. When the matching conditions in the PCC rule are met, the LRF rule is activated.

If you are using Junos OS Subscriber Aware, specify the name of the LRF rule at the **[edit unified-edge pcef pcc-action-profiles *profile-name*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the LRF rule at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level.

Options

lrf-rule-name—LRF rule name. The referenced LRF rule must be configured in an LRF profile.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Activation of an LRF Rule by a PCC Rule](#) | 468

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management | 459](#)

[Configuring an LRF Profile for Subscribers | 460](#)

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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 Policers Based on Logical Interface Bandwidth

[bandwidth-percent | 747](#) statement

[interface-specific](#) statement

logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces)

Syntax

```
logical-interface-fpc-redundancy;
```

Hierarchy Level

```
[edit interfaces aenumber aggregated-ether-options]
```

Release Information

Statement introduced in Junos OS Release 11.2.

Statement introduced in Junos OS Release 13.2R2 for EX Series switches.

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

Configuring Module Redundancy for a Virtual Chassis

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.

Support for PTX series routers with third-generation FPCs added in Junos OS Release 18.3R1.

Description

Configure a logical interface policer. For PTX series routers running Junos OS Release 18.3R1 or later, you can use this command to configure separate firewall filters for different family address types (IPv4 and IPv6) that share the same interface, and configure the same policer as an action for the filter.

To configure the aggregate policer, configure the firewall policer you want to use as **logical-interface-policer**. And at the **firewall family *family-name* filter *filter-name*** hierarchy level where you will reference the policer, make the policer an **interface-specific** firewall filter action.

The sample configuration shows the relationship.

```
firewall {
  policer Shared_Policer {
    logical-interface-policer;
    if-exceeding {
```

```

        bandwidth-limit 100m;
        burst-size-limit 500k;
    }
    then {
        discard;
    }
}
}

```

```

family inet {
    filter filter_name{
        interface-specific;
        term term_name {
            then {
                policer Shared_Policer;
                count cinet;
            }
        }
    }
}
}

```

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

<i>Two-Color and Three-Color Logical Interface Policers</i>
<i>Traffic Policer Types</i>
<i>Configuring and Applying Tricolor Marking Policers</i>
action 707
<i>Configuring Gigabit Ethernet Two-Color and Tricolor Policers</i>
<i>action</i>

logical-system (Subscriber Secure Policy)

Syntax

```
logical-system logical-system-name;
```

Hierarchy Level

```
[edit services radius-flow-tap]
```

Release Information

Statement introduced in Junos OS Release 15.1R3 for enhanced subscriber management on MX Series routers.

Description

Specify the logical system that is used to send mirrored packets to a mediation device for subscriber secure policy traffic mirroring. When you specify a logical system, you must also specify a routing instance.

Options

logical-system-name—Name of the logical system.

Default: Logical system **default**

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Subscriber Secure Policy Overview | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

[Configuring Support for Subscriber Secure Policy Mirroring | 574](#)

login

Syntax

```
login {
  announcement text;
  class class-name {
    allow-hidden-commands;
    no-hidden-commands {
      except ["regular expression or command 1" "regular expression or command 2" ...];
    }
    access-end hh:mm;
    access-start hh:mm;
    ( allow-commands "(regular-expression1)|(regular-expression2)..." | allow-commands-regexps ["regular expression 1" "regular expression 2" ... ] );
    ( allow-configuration "(regular-expression1)|(regular-expression2)..." | allow-configuration-regexps ["regular expression 1" "regular expression 2" ... ] );
    allow-sources [ source-addresses ... ];
    allow-times [ times ... ];
    allowed-days [ days of the week ];
    cli {
      prompt prompt;
    }
    configuration-breadcrumbs;
    confirm-commands ["regular expression or command 1" "regular expression or command 2" ...] {
      confirmation-message;
    }
    ( deny-commands "(regular-expression1)|(regular-expression2)..." | deny-commands-regexps ["regular expression 1" "regular expression 2" ... ] );
    ( deny-configuration "(regular-expression1)|(regular-expression2)..." | deny-configuration-regexps ["regular expression 1" "regular expression 2" ... ] );
    deny-sources [ source-addresses ... ];
    deny-times [ times ... ];
    idle-timeout minutes;
    logical-system logical-system-name;
    login-alarms;
    login-script login-script;
    login-tip;
    no-scp-server;
    no-sftp-server;
    permissions [ permissions ];
    satellite all;
    security-role (audit-administrator | crypto-administrator | ids-administrator | security-administrator);
    tenant tenant-system-name;
  }
}
```

```

deny-sources {
    address [ source-addresses ... ];
}
idle-timeout minutes;
message text;
password {
    change-type (character-sets | set-transitions);
    format (sha1 | sha2 | sha256 | sha512);
    maximum-length length;
    maximum-lifetime days
    minimum-changes number;
    minimum-character-changes number
    minimum-length length;
    minimum-lifetime days
    minimum-lower-cases number;
    minimum-nerics number;
    minimum-punctuations number;
    minimum-reuse number;
    minimum-upper-cases number;
}
retry-options {
    backoff-factor seconds;
    backoff-threshold number;
    lockout-period minutes;
    maximum-time seconds;
    minimum-time seconds;
    tries-before-disconnect number;
}

```

```

user username {
  authentication {
    encrypted-password encrypted-password;
    no-public-keys;
    ssh-eccdsa name {
      from from;
    }
    ssh-ed25519 name {
      from from;
    }
    ssh-rsa name {
      from from;
    }
  }
  cli {
    prompt prompt;
  }
  class class-name;
  full-name full-name;
  uid uid-value;
}

```

Hierarchy Level

[edit system]

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

deny-sources option introduced in Junos OS Release 11.2.

All of the statements and options introduced previously were introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure user access to the device.

Options

announcement text—Configure a system login announcement. This announcement appears after a user logs in. Sometimes you want to make announcements to authorized users only after they have logged in. For example, you might want to announce an upcoming maintenance event.

To display a message before the user logs in, configure a system login message using the **message** statement rather than configuring a system login announcement.

You can format the announcement using the following special characters:

- \n—New line
- \t—Horizontal tab
- \'—Single quotation mark
- \"—Double quotation mark
- \\—Backslash

If the text of the announcement contains any spaces, enclose the text in quotation marks.

Default: No login announcement is displayed.

deny-sources—(Mandatory) Never allow access from these hosts. The source addresses can be IPv4 or IPv6 addresses, prefix lengths, or hostnames.

Syntax: address [source-addresses]

idle-timeout minutes— For a login class, configure the maximum time in minutes that a session can be idle before the session times out and the user is logged out of the device. The session times out after remaining at the CLI operational mode prompt for the specified time.

NOTE: After the user logs in to a device from a shell prompt such as `csh`, if the user starts another program to run in the foreground of the CLI, the idle-timer control is stopped from being computed. The calculation of the idle time of the CLI session is restarted only after the foreground process exits and the control is returned to the shell prompt. When the restart of the idle-timer control occurs, if no interaction from the user occurs on the shell, the user is automatically logged out after the time set on this statement.

Default: If you omit this statement, a user is never forced off the system after extended idle times.

Range: Range: 0 through 4294967295 minutes

NOTE: The idle-timeout feature is disabled if the value of *minutes* is set to 0.

message text—Configure a system login message. A login message displays a banner to users when they access the device, before they log in. To display a message only after the user logs in, configure a system login announcement using the **announcement** statement instead of configuring a system login message.

Before you create any user accounts, it's a good idea to configure an initial login message.

You can format the message using the following special characters:

- \n—New line
- \t—Horizontal tab
- \'—Single quotation mark
- \"—Double quotation mark
- \\—Backslash

If the text of the message contains any spaces, enclose the text in quotation marks.

Default: No login message is displayed.

The remaining statements are explained separately. See [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Defining Junos OS Login Classes

Configuring Junos OS to Display a System Login Announcement

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](#) | 40

[Configuring Schedulers in a Dynamic Profile for Subscriber Access](#) | 52

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 | 707

max-queues-per-interface

Syntax

```
max-queues-per-interface (8 | 4);
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number],  
[edit chassis lcc number fpc slot-number pic pic-number] (Routing Matrix)
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support for TX Matrix and TX Matrix Plus added in Junos OS Release 9.6.

Description

On IQ, MPC, and DPC interfaces on M120, T320, T640, T1600, TX Matrix, and TX Matrix Plus routers, or on MIC or MPC interfaces on MX Series routers, set the number of egress queues per port to four or eight.

NOTE: If you include the **max-queues-per-interface 8** statement, the configuration at the **[edit class-of-service]** hierarchy level must also support eight queues per interface.

NOTE: When you include the **max-queues-per-interface** statement and commit the configuration, all physical interfaces on the PIC/MIC are deleted and readdded. Also, the PIC/MIC is restarted automatically. You should change modes between four queues and eight queues only when there is no active traffic going to the PIC/MIC.

By default, IQ PICs on T Series and M320 routers are restricted to a maximum of four egress queues per interface. If you include the **max-queues-per-interface 4** statement, you can configure all four ports and configure up to four queues per port.

For Quad T3 and Quad E3 PICs and for 4-port OC3c/STM1 Type I and Type II PICs on M320 and T Series routers, when you include the **max-queues-per-interface 8** statement, you can configure up to eight queues on ports 0 and 2. After you commit the configuration, the PIC goes offline and comes back online with only ports 0 and 2 operational. No interfaces can be configured on ports 1 and 3.

NOTE: Starting from Junos OS Release 14.1R8, 14.2R6, 15.1F6, 15.1R3, 15.1R4, and 16.1R1, the restricted queue PICs without the **max-queues-per-interface** configuration boot up with a maximum of eight queues per port and two operational ports (port 0 and 2). PICs with restricted queues include Quad T3 PIC, Quad E3 PIC, 4-port SONET/SDH OC3c/STM1 PIC, and 4-Port OC3 and 1-port OC12 PICs with SFP.

On certain older MPCs (MPC1 Q, MPC1E Q, MPC2 Q, MPC2E Q), you can include the **max-queues-per-interface** statement to set the number of queues per logical interface to four or eight. Setting **max-queues-per-interface 4** sets the MPC to have four queues per logical interface and provides twice as many logical interfaces on the MPC as setting **max-queues-per-interface 8**.

NOTE: For consistency, **max-queues-per-interface** should not be set on MPCs starting from Junos OS 14.1X51.

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 Support Eight Queues on IQ Interfaces for T Series and M320 Routers

Configuring Up to 16 Custom Forwarding Classes

Enabling Eight Queues on ATM Interfaces

[Configuring the Maximum Number of Queues for Trio MPC/MIC Interfaces | 104](#)

Example: Configuring CoS on SRX5000 Devices with an MPC

Example: Enabling Eight-Queue Class of Service on Redundant Ethernet Interfaces on SRX Series Devices in a Chassis Cluster

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 334](#) 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 | 328](#)

[Configuring the Match Order for Fast Update Filters | 333](#)

[Fast Update Filter Match Conditions | 334](#)

maximum-bit-rate (PCC Action Profiles)

Syntax

```
maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],  
[edit services pcef pcc-action-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the maximum bit rate (MBR) that you want a PCC action profile to use for uplink and downlink traffic.

If you are using Junos OS Subscriber Aware, specify the MBR at the **[edit unified-edge pcef pcc-action-profiles *profile-name*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the MBR at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level.

Default

If you configure the **maximum-bit-rate** statement but do not specify MBR values for **uplink** and **downlink**, the default value is 0.

Options

mbr-uplink-value—MBR value for the uplink direction.

Range: 1 through 6144000 Kbps.

mbr-downlink-value—MBR value for the downlink direction.

Range: 1 through 6144000 Kbps.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

member (Application Identification)

Syntax

```
[member member-name];
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type signature I4-I7-signature-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Define a member name for a custom application definition. Custom definitions can contain multiple members that define attributes for an application. You can define a maximum of four member names.

Options

member-name—Name of a member for a custom application definition. You can define a maximum of four member names.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

[Application Identification Overview | 431](#)[Application Identification Overview | 431](#)

mld (Dynamic Profiles)

Syntax

```
mld {
  interface interface-name {
    (accounting | no-accounting);
    disable;
    group-limit limit;
    group-policy;
    group-threshold value;
    immediate-leave;
    log-interval seconds;
    oif-map;
    passive;
    ssm-map ssm-map-name;
    ssm-map-policy ssm-map-policy-name;
    static {
      group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
          source-count number;
          source-increment increment;
        }
      }
    }
    version version;
  }
}
```

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.

You can configure MLD in a dynamic client profile or a dynamic service profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Dynamic MLD Configuration Overview](#) | 401

Enabling MLD

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 in a dynamic client profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring the Multicast Forwarding Cache

Example: Configuring a Multicast Flow Map

Example: Configuring Source-Specific Multicast Groups with Any-Source Override

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 | 553](#)

[Subscriber Secure Policy Support for IPv4 Multicast Traffic | 621](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

[Configuring DTCP-Initiated Subscriber Secure Policy Mirroring Overview | 590](#)

netconf (Remote Device Management)

Syntax

```
netconf {
  bulk-interval milliseconds;
  bulk-limit number;
  connection-retry-interval seconds;
  password password;
  port port-number;
  reconfigure-bulk-limit number;
  response-timeout seconds;
  response-timeout-count number;
  user-name name;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name system services remote-device-management service-device device-name
  provisioning-method],
[edit system services remote-device-management service-device device-name provisioning-method]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Configure the NETCONF XML management protocol as the method for provisioning and deprovisioning services on the remote device.

Options

bulk-interval *milliseconds*—(Optional) Interval during which multiple services are provisioned or deprovisioned based on the assigned dictionary before the configuration is committed to the service device. When the interval times out, the service actions are committed in bulk before additional actions for the device can take place. You can use the interval (together with the **bulk-limit** option) to optimize your service device configuration during scaled subscriber negotiation and service provisioning or subscriber termination and service deprovisioning.

NOTE: The **bulk-interval** configuration is ignored when the **bulk-limit** is set to 1.

Range: 500 through 5000

Default: 1

bulk-limit *number*—(Optional) Maximum number of services provisioned or deprovisioned based on the assigned dictionary during the bulk interval before the configuration is committed to the service device. When the limit is reached, the service actions are committed in bulk before additional actions for the device can take place. You can use the limit (together with the **bulk-interval** option) to optimize your service device configuration during scaled subscriber negotiation and service provisioning or subscriber termination and service deprovisioning.

Range: 1 through 1000

Default: 1

connection-retry-interval *seconds*—(Optional) The interval between successive attempts to establish a NETCONF session with the remote device.

Range: 1 through 10

Default: 3

password *password*—Password used by the NETCONF protocol to access the remote device during service management. The maximum length of the password is 64 bytes.

NOTE: If you change the password when any active subscriber services are mapped to the device, the change takes effect only when the device reconnects.

port *port-number*—(Optional) TCP port number for the NETCONF session.

NOTE: You cannot change the port number for the device when any active subscriber services are mapped to it.

Range: 1 through 65,535

Default: 830 (NETCONF over SSH)

reconfigure-bulk-limit *number*—(Optional) When the device is reconfigured, this is the maximum number of services provisioned or deprovisioned on the service device for the access domain based on the assigned dictionary before the configuration is committed to the service device. When the limit is reached, the service actions are committed in bulk before additional actions for the device can take place.

If the access domain is not yet configured, this option has no effect, because there are no matching services to install.

Range: 1 through 1000

Default: 100

response-timeout *seconds*—(Optional) Period during which the device must respond to an attempt to provision or deprovision a service. For a subscriber service over a NETCONF connection, the timeout is a failure equivalent to an explicit failure response received from the device.

Range: 1 through 180

Default: 3

response-timeout-count *number*—(Optional) Number of consecutive response timeouts that can occur before the BNG takes action. The default action is to close and reopen the NETCONF connection.

Range: 1 through 10

Default: 3

user-name *name*—Name used to access the remote device during service management. The maximum length of the name is 64 bytes.

NOTE: If you change the username when any active subscriber services are mapped to the device, the change takes effect only when the device reconnects.

Required Privilege Level

system—To view this statement in the configuration.

RELATED DOCUMENTATION

[Configuring Remote Device Management for Service Provisioning | 648](#)

[Remote Device Services Manager \(RDSM\) Overview | 631](#)

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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Recording IGMP Join and Leave Events

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*

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.

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

[Dynamic MLD Configuration Overview](#) | 401

Example: Configuring Multicast with Subscriber VLANs

order (Application Identification)

Syntax

```
order order;
```

Hierarchy Level

```
[edit services application-identification application name address-mapping name],
[edit services application-identification application application-name icmp-mapping],
[edit services application-identification application application-name ip-protocol-mapping],
[edit services application-identification application application-name over protocol-type signature I4-I7-signature-name],
[edit services application-identification application application-name
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Support at the **[edit services application-identification application *application-name*]** hierarchy level introduced in Junos OS Release 19.3R1 on MX Series running Next Gen Services.

Description

Define application matching priority. For address configurations, the order number resolves the conflict when multiple address entries are matched for a specific session. The lower number has a higher priority.

Options

order—Order sequence number. This value is mandatory and must be unique.

Default: 0

Range: 0 through 65,535

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

[Application Identification Overview | 431](#)[Application Identification Overview | 431](#)

order-priority (Application Identification)

Syntax

```
order-priority (high | low);
```

Hierarchy Level

```
[edit services application-identification application application-name address-mapping name],  
[edit services application-identification application application-name icmp-mapping],  
[edit services application-identification application application-name ip-protocol-mapping],  
[edit services application-identification application application-name over protocol-type signature I4-I7-signature-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Define the priority of signatures when both a custom signature and predefined signature apply to a protocol bundle.

Options

high—Custom signatures have priority over predefined signatures.

low—Predefined signatures have priority over custom signatures.

Default: high

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

[Application Identification Overview | 431](#)[Application Identification Overview | 431](#)

output (Dynamic Service Sets)

Syntax

```
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* [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.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Options

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](#) | 365

[Associating Service Sets with Interfaces in a Dynamic Profile](#) | 366

outputs (Analytics)

Syntax

```
outputs {  
  file {  
    parameters {  
      path file-path;  
    }  
  }  
  kafka {  
    parameters {  
      server ip-address;  
      topic topic-name;  
      encoding encoding-type;  
    }  
  }  
  output-ipfix {  
    parameters {  
      collector-address ip-address;  
      collector-ca-certificate file-path;  
      collector-certificate file-path;  
      collector-certificate-key file-path;  
      collector-connection-retry-interval seconds;  
      collector-tcp-port port-number;  
      collector-vrf-name vrf-name;  
    }  
  }  
}
```

Hierarchy Level

[edit services analytics agent [service-agents](#) *agent-name*]

Release Information


Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Statement introduced in Junos OS Release 18.4R1 on PTX Series routers.

kafka and **file** options added in Junos OS Release 18.4R1 on MX Series and PTX Series routers.

Description

Configure parameters for the Network Telemetry Framework (NTF) agent output plug-in.



NOTE: When you modify the output plug-in configuration of a service agent instance, the associated service agent daemon is restarted.

Options

file parameters—Configure parameters for sending data in a log file to a data collector.

path *pathname*—Path for the log file to which to save the data. For example, **path /tmp/example_file.log**

kafka parameters—Configure parameters for sending data to a Kafka data collector.

server *ip-address*—IP address of the Kafka server.

topic *filename*—Kafka topic name. The naming convention of the topic is *server-name.jti.encoding-type*. The encoding type options are **avro**, **json**, or **msgpack**.

encoding *encoding-type*—Encoding type. Options are **avro**, **json**, or **msgpack**.

output-ipfix parameters—Configure parameters for the IPFIX mediation service agent to send the IPFIX records that have been consolidated on the router to the IPFIX collector.

You must configure the IP address of the upstream IPFIX collector. When you optionally configure at least one of the collector certificate options (**collector-ca-certificate**, **collector-certificate**, and **collector-certificate-key**), the IPFIX mediator attempts to use TLS to connect with the collector. Otherwise, the mediator uses a TCP connection.

NOTE: Any change you make to an existing **output-ipfix** output plug-in configuration restarts the IPFIX service agent daemon to apply the changes.

collector-address *ip-address*—IP address of the upstream IPFIX collector.

collector-ca-certificate *file-path*—(Optional) Path for the certificate, provided by a trusted certificate authority (CA), that is used to sign the peer certificate at the peer (IPFIX collector) level. The certificate is expected to be in .pem container format.

collector-certificate *file-path*—(Optional) Path for the client certificate that the server (IPFIX collector) uses to authenticate the client and enable mutual authentication. The fully-qualified domain name (FQDN) of both the client and the server are stored in the certificate's Subject Alternative Name field when the client and server certificates are generated. The certificate is expected to be in .pem container format.

collector-certificate-key *file-path*—(Optional) Private key file that is loaded to decrypt the encrypted message sent from the peer.

collector-connection-retry-interval *seconds*—(Optional) Interval in seconds at which the output plug-in retries connecting to the IPFIX collector.

Range: 1 through 25

Default: 20

collector-tcp-port *port-number*—(Optional) Number of the TCP port used to connect to the IPFIX collector.

Default: 4740

collector-vrf-name *vrf-name*—(Optional) Name of the VRF (routing instance) in which IPFIX packets are routed.

Default: default

Required Privilege Level

system

RELATED DOCUMENTATION

[Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data | 672](#)

[Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator | 680](#)

Configuring NTF Agent

[IPFIX Mediation on the BNG | 667](#)

[Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector | 677](#)

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 | 40](#)

[Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile | 217](#)

[Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers | 50](#)

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[CoS Adjustment Control Profiles Overview | 183](#)

[Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates | 111](#)

[Bandwidth Management for Downstream Traffic in Edge Networks Overview | 109](#)

[egress-shaping-overhead](#)

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—Enables IGMP to receive control traffic on the interface.

send-general-query—Enables IGMP to send general queries on the interface.

send-group-query—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 User Guide*.

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

[Dynamic MLD Configuration Overview](#) | 401

Example: Configuring Multicast with Subscriber VLANs

path (Steering)

Syntax

```
path {
  ipv4-address ipv4-address;
  ipv6-address ipv6-address;
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name steering],
[edit services pcef pcc-action-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the IP address of a third-party server to which the PCC action profile steers HTTP traffic for applying services. If you configure this, the PCC action profile can only be used in PCC rules that match only HTTP-based applications and all flows.

The remaining statements are explained separately.

Required Privilege Level

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

[Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment](#)

pattern (Application Identification)

Syntax

```
pattern pattern;
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type signature l4-l7-signature-name  
  member member-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Define an attack pattern to be detected.

Options

pattern—User-defined pattern of attack to match, using a regular expression.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

[Application Identification Overview | 431](#)[Application Identification Overview | 431](#)

pcc-action-profile (PCC Rules)

Syntax

```
pcc-action-profile profile-name;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rules-name then],  
[edit services pcef pcc-rules rules-name then]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-rules *rules-name* then]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the name of the action profile to include in a policy and charging control (PCC) rule configuration. The action profile defines the treatment to be applied to specific service data flows or to packets associated with specific applications.

If you are using Junos OS Subscriber Aware, specify the name of the action profile at the **[edit unified-edge pcef pcc-rules *rules-name* then]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the action profile at the **[edit services pcef pcc-rules *rules-name* then]** hierarchy level.

Options

profile-name—Name of the PCC action profile that the PCC rule references. The referenced action profile must be configured.

Range: 1 through 63 characters.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Policy and Charging Control Rules | 418](#)

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

pcc-action-profiles

Syntax

```
pcc-action-profiles profile-name {
  forwarding-class class-name;
  gate-status (uplink | downlink | uplink-downlink | disable-both);
  hcm-profile hcm-profile-name;
  logging-rule lrf-rule-name;
  maximum-bit-rate uplink mbr-uplink-value downlink mbr-downlink-value;
  monitoring-key key_string;
  redirect {
    url url-name;
  }
  steering {
    keep-existing-steering;
    path {
      ipv4-address ipv4-address;
      ipv6-address ipv6-address;
    }
    routing-instance {
      downlink downlink-vrf-name;
      uplink uplink-vrf-name;
    }
  }
}
```

Hierarchy Level

```
[edit unified-edge pcef],
[edit services pcef]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Configure a PCC action profile. A PCC action profile defines the treatment to be applied to specific service data flows or to packets associated with specific applications. A PCC action profile is specified in the **then** clause of a PCC rule.

If you are using Junos OS Subscriber Aware, configure the PCC action profile at the **[edit unified-edge pcef]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the PCC action profile at the **[edit services pcef]** hierarchy level. The following options are not applicable to subscriber management:

- **hcm-profile**
- **monitoring-key**

Options

profile-name—Name of the PCC action profile.

Range: 1 through 63 characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

[Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment](#)

[Configuring TDF Subscriber Usage Monitoring for Traffic That Matches Predefined PCC Rules](#)

[Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware](#)

pcc-context

Syntax

```
pcc-context {
  input-service-filter-name filter-name;
  input-service-set-name service-set-name;
  ipv6-input-service-filter-name filter-name;
  ipv6-input-service-set-name service-set-name;
  ipv6-output-service-filter-name filter-name;
  ipv6-output-service-set-name service-set-name;
  output-service-filter-name filter-name;
  output-service-set-name service-set-name;
  profile-name pcef-profile-name;
}
```

Hierarchy Level

[edit access [profile](#) *profile-name* [session-options](#)]

Release Information

Statement introduced in Junos OS Release 18.2R1 on MX Series.

Description

Specify the PCEF profile that contains the policy and charging control (PCC) rules that a PCRF can directly activate for a subscriber, and specify the input and output service sets that process the PCC rules. You can optionally specify service filters for the service sets. PCC rules specify the policies to apply to traffic based on the application being used by the subscriber or the Layer 3 and Layer 4 service data flow information.

Options

input-service-filter-name *filter-name*—(Optional) Input IPv4 service filter name. The service filter identifies conditions for which you want to skip application-aware policy control.

input-service-set-name *service-set-name*—Input IPv4 service-set name. Use a service set that is enabled for application-aware policy control.

ipv6-input-service-filter-name *filter-name*—(Optional) Input IPv6 service filter name. The service filter identifies conditions for which you want to skip application-aware policy control.

ipv6-input-service-set-name *service-set-name*—Input IPv6 service set name. Use a service set that is enabled for application-aware policy control.

ipv6-output-service-filter-name *filter-name*—(Optional) Output IPv6 service filter name. The service filter identifies conditions for which you want to skip application-aware policy control.

ipv6-output-service-set-name *service-set-name*—Output IPv6 service set name. Use a service set that is enabled for application-aware policy control.

output-service-filter-name *filter-name*— (Optional) Output IPv4 service filter name. The service filter identifies conditions for which you want to skip application-aware policy control.

output-service-set-name *service-set-name*—Output IPv4 service-set name. Use a service set that is enabled for application-aware policy control.

profile-name *pcef-profile-name*—PCEF profile name. The PCEF profile contains PCC rules or PCC rulesets that the PCRF can directly activate. The PCEF profile must be configured under **dynamic-policy-control** at the `[edit services pcef profiles profile-name]` hierarchy level.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management](#) | 428

pcc-rule

Syntax

```
[pcc-rule rule-name precedence number];
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rule-bases rulebase-name],  
[edit services pcef pcc-rule-bases rulebase-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-rule-bases *rulebase-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Specify one or more policy and charging control (PCC) rules and the rules precedence in a PCC rulebase.

If you are using Junos OS Subscriber Aware, configure the PCC rules at the **[edit unified-edge pcef pcc-rule-bases *rulebase-name*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the PCC rules at the **[edit services pcef pcc-rule-bases *rulebase-name*]** hierarchy level.

Options

rule-name—Name of the PCC rule. The referenced PCC rule must be configured.

Range: 1 through 63 characters.

number—Precedence value assigned to the PCC rule. The precedence assigned must be unique among the configured PCC rules.

Range: 1 through 65,535

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring a Policy and Charging Control Rulebase | 421

Configuring Policy and Charging Control Rules | 418

pcc-rulebases (PCEF)

Syntax

```
pcc-rulebases rulebase-name {
    [pcc-rule rule-name precedence number];
}
```

Hierarchy Level

```
[edit unified-edge pcef],
[edit services pcef]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Configure a policy and charging control (PCC) rulebase. You can specify from 1 through 4000 rules in a rulebase.

If you are using Junos OS Subscriber Aware, configure the PCC rulebase at the **[edit unified-edge pcef]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the PCC rulebase at the **[edit services pcef]** hierarchy level.

Options

rulebase-name—Name of the PCC rulebase.

Range: 1 through 63 characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring a Policy and Charging Control Rulebase](#) | 421

pcc-rulebases (PCEF Profile)

Syntax

```
[pcc-rulebases rulebase-name <time-of-day-profile profile-name>];
```

Hierarchy Level

```
[edit unified-edge pcef profiles profile-name aaa-policy-control],
[edit unified-edge pcef profiles profile-name dynamic-policy-control],
[edit unified-edge pcef profiles profile-name static-policy-control],
[edit services pcef profiles profile-name static-policy-control],
[edit services pcef profiles profile-name dynamic-policy-control]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef profiles *profile-name* static-policy-control]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support at the **[edit services pcef profiles *profile-name* dynamic-policy-control]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 18.2R1 on MX Series.

Description

Specify a policy and charging control (PCC) rulebase for a policy control profile.

If you are using Junos OS Subscriber Aware, specify the PCC rulebase at the **[edit unified-edge pcef profiles *profile-name* (aaa-policy-control | dynamic-policy-control | static-policy-control)]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the PCC rulebase at the **[edit services pcef profiles *profile-name* (static-policy-control | dynamic-policy-control)]** hierarchy level.

Options

rulebase-name—Name of the PCC rulebase. The referenced PCC rulebase must be configured.

time-of-day-profile profile-name—(Optional; only applies to rulebases in static PCEF profiles for Junos OS Subscriber Aware) Use the specified name of the time-of-day profile to apply to the PCC rulebase. The referenced profile must already be defined at the **[edit unified-edge pcef]** hierarchy level. The time-of-day profile specifies the time of day, day of the week, or day of the month to activate or deactivate the PCC rulebase for subscribers assigned to the PCEF profile.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Static Policies

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls

Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile

[Configuring a Policy and Charging Control Rulebase | 421](#)

[Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management | 423](#)

pcc-rules (PCEF)

Syntax

```
pcc-rules rule-name {
  from {
    <application-groups [application-group-name]>;
    <applications [application-name]>;
    flows ([flow-identifier | any]);
  }
  then {
    pcc-action-profile profile-name;
  }
}
```

Hierarchy Level

```
[edit unified-edge pcef],
[edit services pcef]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Configure the PCC rules. A PCC rule identifies the subscriber IP packets that are associated with a service data flow (SDF) or application and defines the treatment to be applied to the packets.

If you are using Junos OS Subscriber Aware, configure the PCC rule at the **[edit unified-edge pcef]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the PCC rule at the **[edit services pcef]** hierarchy level.

Options

rule-name—Name of the PCC rule.

Range: 1 through 63 characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Policy and Charging Control Rules | 418](#)

Configuring TDF Subscriber Usage Monitoring for Traffic That Matches Predefined PCC Rules

pcc-rules (PCEF Profile)

Syntax

```
pcc-rules [rule-name precedence number <time-of-day-profile profile-name>];
```

Hierarchy Level

```
[edit unified-edge pcef profiles profile-name aaa-policy-control],
[edit unified-edge pcef profiles profile-name dynamic-policy-control],
[edit unified-edge pcef profiles profile-name static-policy-control],
[edit services pcef profiles profile-name static-policy-control],
[edit services pcef profiles profile-name dynamic-policy-control]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef profiles *profile-name* static-policy-control]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support at the **[edit services pcef profiles *profile-name* dynamic-policy-control]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 18.2R1 on MX Series.

Description

Specify the policy and charging control (PCC) rules for a policy and charging enforcement function (PCEF) profile and assign a precedence to each PCC rule. You can configure up to 32 PCC rules in a PCEF profile.

If you are using Junos OS Subscriber Aware, specify the PCC rules at the **[edit unified-edge pcef profiles *profile-name* (aaa-policy-control | dynamic-policy-control | static-policy-control)]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the PCC rules at the **[edit services pcef profiles *profile-name* (static-policy-control | dynamic-policy-control)]** hierarchy level.

Options

rule-name—Name of the PCC rule. The referenced PCC rule must be configured.

precedence *number*—Use the specified precedence value assigned to a PCC rule. A lower precedence value indicates a higher precedence.

Range: 1 through 65,535

time-of-day-profile *profile-name*—(Optional; only applies to rules in static PCEF profiles for Junos OS Subscriber Aware) Use the specified name of the time-of-day profile to apply to the PCC rule. The referenced profile must already be defined at the **[edit unified-edge pcef]** hierarchy level. The time-of-day profile specifies the time of day, day of the week, or day of the month to activate or deactivate the PCC rule for subscribers assigned to the PCEF profile.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Static Policies

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls

Configuring Static Time-of-Day PCC Rule Activation and Deactivation in a Junos OS Subscriber Aware PCEF Profile

[Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management | 423](#)

[Configuring Policy and Charging Control Rules | 418](#)

pcef (Dynamic Profiles)

Syntax

```
pcef pcef-profile-name {
    activate rule-name | activate-all;
}
```

Hierarchy Level

```
[edit dynamic-profile profile-name interfaces interface-name unit logical-unit-number service]
```

Release Information

Statement introduced in Junos OS Release 17.2 on MX Series.

Description

Assign a policy and charging enforcement function (PCEF profile) to the dynamic profile. The PCEF profile specifies a set of PCC rules and rulebases to assign to a subscriber for application-aware policy control, and assigns a precedence value to each predefined rule. Also specify which of the PCC rules to activate for the dynamic profile.

Options

activate rule-name—Name of the PCC rule to activate or the predefined dynamic interface variable **\$junos-pcef-rule**. To specify more than one rules, include this line multiple times.

activate-all—Activates all of the PCC rules.

pcef-profile-name—Name of the PCEF profile or the predefined dynamic interface variable **\$junos-pcef-profile**.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Application-Aware Policy Control for Subscriber Management](#) | 410

pcef-profile (Service Set)

Syntax

```
pcef-profile pcef-profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Specify the dummy PCEF profile that you configured at the **[edit services pcef]** hierarchy level. This profile is a placeholder profile with no configuration options, but it must be specified to enable PCEF functionality on the services plane.

Options

pcef-profile-name—Name of the PCEF profile.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Applying Services to Subscriber-Aware Traffic with a Service Set

[Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control | 424](#)

[Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management | 423](#)

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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

<i>Three-Color Policer Configuration Overview</i>
<i>Policer Bandwidth and Burst-Size Limits</i>
<i>Policer Color-Marking and Actions</i>
<i>Dual Token Bucket Algorithms</i>
<i>Determining Proper Burst Size for Traffic Policers</i>
committed-burst-size 788
committed-information-rate 790
excess-burst-size 847
peak-information-rate 1032

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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 on EX, M, and T Series routers
- 1500 through 18,446,744,073,709,551,615 bps on Mx Series routers

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Three-Color Policer Configuration Overview</i>
<i>Policer Bandwidth and Burst-Size Limits</i>
<i>Policer Color-Marking and Actions</i>
<i>Dual Token Bucket Algorithms</i>
<i>Determining Proper Burst Size for Traffic Policers</i>
committed-burst-size 788
committed-information-rate 790
excess-burst-size 847
peak-burst-size 1030

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for PTX series routers with third-generation FPCs added in Junos OS Release 18.3R1.

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 in aggregate 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

	<i>Two-Color and Three-Color Physical Interface Policers</i>
	<i>physical-interface-filter</i>

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.

You can configure the policer in static firewall filters or dynamic firewall filters in a dynamic client profile or a dynamic service profile.

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**.

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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

Single-Rate Two-Color Policer Overview

Using Multifield Classifiers to Set Packet Loss Priority

[filter \(Configuring\)](#) | **866**

priority (Schedulers)

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

policy-based-logging (LRF Profile)

Syntax

```
policy-based-logging;
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure policy-based logging, which causes the LRF rules to be activated by PCC rules.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management | 459](#)

policy-options (Dynamic Profiles)

Syntax

```
policy-options {
  prefix-list uid {
    ip-addresses;
    dynamic-db;
  }
}
```

Hierarchy Level

[edit **dynamic-profiles** *profile-name*]

Release Information

Statement introduced before Junos OS Release 11.4.

Description

Define a list of IPv4 or IPv6 address prefixes for use in a dynamic firewall filter or in an HTTP redirect configuration.

You can configure up to 85,325 prefixes in each prefix list. To configure more than 85,325 prefixes, configure multiple prefix lists and apply them to multiple firewall filter terms.

You can configure policy options in a dynamic client profile or a dynamic service profile.

Options

uid—Unique identifier of the prefix list. You must assign a UID as the prefix list name.

ip-addresses—List of IPv4 or IPv6 address prefixes, one IP address per line in the configuration.

dynamic-db—Specify that the routing policy and policy objects reference policies configured in the dynamic database at the **[edit dynamic]** hierarchy level.

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*

policy-statement

Syntax

```

policy-statement policy-name {
  term term-name {
    from {
      as-path-unique-count count (equal | orhigher | orlower);
      family family-name;
      match-conditions;
      policy subroutine-policy-name;
      prefix-list prefix-list-name;
      prefix-list-filter prefix-list-name match-type <actions>;
      protocol protocol-name;
      route-filter destination-prefix match-type <actions>;
      source-address-filter source-prefix match-type <actions>;
      tag value;
      traffic-engineering;
    }
    to {
      match-conditions;
      policy subroutine-policy-name;
    }
    then actions;
  }
  then {
    aggregate-bandwidth;
    dynamic-tunnel-attributes dynamic-tunnel-attributes;
    limit-bandwidth limit-bandwidth;
    multipath-resolve multipath-resolve;
    no-entropy-label-capability;
    prefix-segment {
      index index;
      node-segment;
    }
    priority (high | medium | low);
    resolution-map map-name;
  }
}

```

Hierarchy Level

```

[edit dynamic-profiles profile-name policy-options],
[edit logical-systems logical-system-name policy-options],

```

[edit policy-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.

inet-mdt option introduced in Junos OS Release 10.0R2.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

route-target option introduced in Junos OS Release 12.2.

Statement introduced in Junos OS 14.1X53-D20 for the OCX Series.

protocol and **traffic-engineering** options introduced in Junos OS Release 14.2.

no-entropy-label-capability option introduced in Junos OS Release 15.1.

priority and **tag value** options introduced in Junos OS Release 17.1.

as-path-unique-count option introduced in Junos OS Release 17.2R1.

prefix-segment option introduced in Junos OS Release 17.2R1 for MX Series routers, PTX Series routers, QFX5100 switches, and QFX10000 switches.

multipath-resolve and **dynamic-tunnel-attributes** options introduced in Junos OS Release 17.3R1.

aggregate-bandwidth and **limit-bandwidth** *limit-bandwidth* options introduced in Junos OS Release 17.4R1 for MX Series, PTX Series, and QFX Series.

l-isis and *l-ospf* keywords at the **protocol** option is introduced in Junos OS Release 19.1R1.

resolution-map statement introduced in Junos OS Release 19.2R1-S1 on MX and PTX Series routers.

lsp and **lsp-regex** options introduced in Junos OS Release 19.4R1.

Description

Define a routing policy, including subroutine policies.

A *term* is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of **accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the **[edit policy-options]** hierarchy level by **policy-statement** *policy-name* in alphabetical order, enter the **show policy-options** configuration command.

The statements are explained separately.

Options

actions—(Optional) One or more actions to take if the conditions match. The actions are described in *Configuring Flow Control Actions*.

family *family-name*—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**. For BGP route target VPN traffic, specify **route-target**. For traffic engineering, specify **traffic-engineering**.

NOTE: When **family** is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the *family* statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

from—(Optional) Match a route based on its source address.

as-path-unique-count *count* (*equal* | *orhigher* | *orlower*)—(Optional) Specify a number from 0 through 1024 to filter routes based on the number of unique autonomous systems (ASs) in the AS path. Specify the match condition for the unique AS path count.

aggregate-bandwidth—(Optional) Enable BGP to advertise aggregate outbound link bandwidth for load balancing.

dynamic-tunnel-attributes *dynamic-tunnel-attributes*—(Optional) Choose a set of defined dynamic tunnel attributes for forwarding traffic over V4oV6 tunnels.

match-conditions—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Routing Policy Match Conditions*.

multipath-resolve *multipath-resolve*—(Optional) Enable the use of all paths for resolution over the specified prefix.

limit-bandwidth *limit-bandwidth*—(Optional) Specify the limit for advertised aggregate outbound link bandwidth for load balancing.

Range: 0 through 4,294,967,295 bytes

no-entropy-label-capability—(Optional) Disable the entropy label capability advertisement at egress or transit routes specified in the policy.

priority (*high* | *medium* | *low*)—(Optional) Configure the priority for an IS-IS route to change the default order in which the routes are installed in the routing table, in the event of a network topology change.

policy subroutine-policy-name—Use another policy as a match condition within this policy. The name identifying the subroutine policy 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 (" "). Policy names cannot take the form `__.*-internal__`, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

policy-name—Name that identifies the policy. 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 (" ").

prefix-list prefix-list-name—Name of a list of IPv4 or IPv6 prefixes.

prefix-list-filter prefix-list-name—Name of a prefix list to evaluate using qualifiers; **match-type** is the type of match, and **actions** is the action to take if the prefixes match.

protocol protocol-name—Name of the protocol used to control traffic engineering database import at the originating point.

Starting in Junos OS Release 19.1R1, you can specify options to match label IS-IS and label OSPF routes using the **l-isis** and **l-ospf** options, respectively. The **isis** options matches all IS-IS routes, excluding labelled IS-IS routes. The **ospf** option matches all OSPF routes, including OSPFv2, OSPFv3 and labelled OSPF routes.

resolution-map—(Optional) Set resolution map modes. A given resolution-map can be shared across multiple policy-statements.

route-filter destination-prefix match-type <actions>—(Optional) List of routes on which to perform an immediate match; **destination-prefix** is the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see *Configuring Route Lists*), and **actions** is the action to take if the **destination-prefix** matches.

source-address-filter source-prefix match-type <actions>—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. **source-prefix** is the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see *Configuring Route Lists*), and **actions** is the action to take if the **source-prefix** matches.

tag value—(Optional) A numeric value that identifies a route. You can tag certain routes to prioritize them over other routes. In the event of a network topology change, Junos OS updates these routes in the routing table before updating other routes with lower priority. You can also tag some routes to identify and reject them based on your requirement.

term term-name—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions.

to—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

then—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

dynamic-db

Understanding Source Packet Routing in Networking (SPRING)

port (LRF Profile)

Syntax

```
port collector-port-number;
```

Hierarchy Level

```
[edit services lrf profile profile-name collector collector-name destination]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Specify the destination port of the collector.

Options

collector-port-number—Port number for the destination address of the collector.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management | 459](#)

port-range (Application Identification)

Syntax

```
port-range {  
    tcp [port-range];  
    udp [port-range];  
}
```

Hierarchy Level

[edit services application-identification application *application-name* over *protocol-type* signature *I4-I7-signature-name*]

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Define TCP or UDP port number range.

Options

port-range—Numeric port ranges. The format for numeric port ranges is in the format *minimum-value-maximum-value*.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

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 post-service filter on the input side of the interface only.

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 | 365](#)

[Associating Service Sets with Interfaces in a Dynamic Profile | 366](#)

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	183
Configuring CoS Adjustment Control Profiles	186
Verifying the CoS Adjustment Control Profile Configuration	188
adjustment-control-profiles	712
application (Adjustment Control Profiles)	729

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 interface-name unit logical-unit-number filter input filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number 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.

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

[Firewall Filters Overview](#)

[Understanding Dynamic Firewall Filters | 226](#)

[Classic Filters Overview | 229](#)

[Fast Update Filters Overview | 323](#)

[Basic Classic Filter Syntax | 232](#)

[Basic Fast Update Filter Syntax | 327](#)

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](#)

[Guidelines for Applying Traffic Policers](#)

Hierarchical Policer Configuration Overview

Hierarchical Policers

[aggregate \(Hierarchical Policer\)](#) | **718**

bandwidth-limit (Hierarchical Policer)

[burst-size-limit \(Hierarchical Policer\)](#) | **752**

[hierarchical-policer](#) | **911**

[if-exceeding \(Hierarchical Policer\)](#) | **920**

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 | 40](#)

[Configuring Schedulers in a Dynamic Profile for Subscriber Access | 52](#)

Dynamic Variables Overview

[scheduler \(Dynamic Scheduler Maps\) | 1134](#)

priority (Application Identification With Next Gen Services)

Syntax

```
priority (high | low);
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R1 on MX Series running Next Gen Services.

Description

Define the priority of signatures when both a custom signature and predefined signature apply to a protocol bundle. This statement is valid only if you are running Next Gen Services. If you are not running Next Gen Services, use the **order-priority** statement.

Options

high—Custom signatures have priority over predefined signatures.

low—Predefined signatures have priority over custom signatures.

Default: high

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

profile (Access)

Syntax

```

profile profile-name {
  accounting {
    address-change-immediate-update
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    ancp-speed-change-immediate-update;
    coa-immediate-update;
    coa-no-override service-class-attribute;
    duplication;
    duplication-filter;
    duplication-vrf {
      access-profile-name profile-name;
      vrf-name vrf-name;
    }
    immediate-update;
    order [ accounting-method ];
    send-acct-status-on-config-change;
    statistics (time | volume-time);
    update-interval minutes;
    wait-for-acct-on-ack;
  }
  accounting-order (radius | [accounting-order-data-list]);
  authentication-order [ authentication-methods ];
  client client-name {
    chap-secret chap-secret;
    group-profile profile-name;
    ike {
      allowed-proxy-pair {
        remote remote-proxy-address local local-proxy-address;
      }
      pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
      ike-policy policy-name;
      interface-id string-value;
    }
    l2tp {
      aaa-access-profile profile-name;
      interface-id interface-id;
      lcp-renegotiation;
      local-chap;
      maximum-sessions number;
      maximum-sessions-per-tunnel number;
    }
  }
}

```

```

    multilink {
        drop-timeout milliseconds;
        fragment-threshold bytes;
    }
    override-result-code session-out-of-resource;
    ppp-authentication (chap | pap);
    ppp-profile profile-name;
    service-profile profile-name(parameter)&profile-name;
    sessions-limit-group limit-group-name;
    shared-secret shared-secret;
}
pap-password pap-password;
ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-ip-address ip-address;
    framed-pool framed-pool;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
}
user-group-profile profile-name;
}
domain-name-server;
domain-name-server-inet;
domain-name-server-inet6;
local {
    flat-file-profile profile-name;
}
preauthentication-order preauthentication-method;
provisioning-order (gx-plus | jsr | pcrf);

```

```

radius {
  accounting-server [ ip-address ];
  attributes {
    exclude {
      attribute-name packet-type;
      standard-attribute number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
      }
      vendor-id id-number {
        vendor-attribute vsa-number {
          packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
        }
      }
    }
  }
  ignore {
    dynamic-iflset-name;
    framed-ip-netmask;
    idle-timeout;
    input-filter;
    logical-system:routing-instance;
    output-filter;
    session-timeout;
    standard-attribute number;
    vendor-id id-number {
      vendor-attribute vsa-number;
    }
  }
}
authentication-server [ ip-address ];
options {
  accounting-session-id-format (decimal | description);
  calling-station-id-delimiter delimiter-character;
  calling-station-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    mac-address;
    nas-identifier;
    stacked-vlan;
    vlan;
  }
  chap-challenge-in-request-authenticator;
  client-accounting-algorithm (direct | round-robin);
}

```

```

client-authentication-algorithm (direct | round-robin);
coa-dynamic-variable-validation;
ethernet-port-type-virtual;
interface-description-format {
    exclude-adapter;
    exclude-channel;
    exclude-sub-interface;
}
juniper-access-line-attributes;
nas-identifier identifier-value;
nas-port-extended-format {
    adapter-width width;
    ae-width width;
    port-width width;
    pw-width width;
    slot-width width;
    stacked-vlan-width width;
    vlan-width width;
    atm {
        adapter-width width;
        port-width width;
        slot-width width;
        vci-width width;
        vpi-width width;
    }
}
nas-port-id-delimiter delimiter-character;
nas-port-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    nas-identifier;
    order {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        interface-text-description;
        nas-identifier;
        postpend-vlan-tags;
    }
    postpend-vlan-tags;
}

```

```

nas-port-type {
    ethernet {
        port-type;
    }
}
override {
    calling-station-id remote-circuit-id;
    nas-ip-address tunnel-client-gateway-address;
    nas-port tunnel-client-nas-port;
    nas-port-type tunnel-client-nas-port-type;
}
remote-circuit-id-delimiter;
remote-circuit-id-fallback {
    remote-circuit-id-format;
    agent-circuit-id;
    agent-remote-id;
}
revert-interval interval;
service-activation {
    dynamic-profile (optional-at-login | required-at-login);
    extensible-service (optional-at-login | required-at-login);
}
vlan-nas-port-stacked-format;
}
preauthentication-server ip-address;
}
radius-server server-address {
    accounting-port port-number;
    accounting-retry number;
    accounting-timeout seconds;
    dynamic-request-port
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    max-outstanding-requests value;
    source-address source-address;
    timeout seconds;
}

```

```

service {
  accounting {
    statistics (time | volume-time);
    update-interval minutes;
  }
  accounting-order (activation-protocol | local | radius);
}
session-limit-per-username number;
session-options {
  client-idle-timeout minutes;
  client-idle-timeout-ingress-only;
  client-session-timeoutminutes;
  pcc-context {
    input-service-filter-name filter-name;
    input-service-set-name service-set-name;
    ipv6-input-service-filter-name filter-name;
    ipv6-input-service-set-name service-set-name;
    ipv6-output-service-filter-name filter-name;
    ipv6-output-service-set-name service-set-name;
    output-service-filter-name filter-name;
    output-service-set-name service-set-name;
    profile-name pcef-profile-name;
  }
  strip-user-name {
    delimiter [ delimiter ];
    parse-direction (left-to-right | right-to-left);
  }
}
subscriber username {
  delegated-pool delegated-pool-name;
  framed-ip-address ipv4-address;
  framed-ipv6-pool ipv6-pool-name;
  framed-pool ipv4-pool-name;
  password password;
  target-logical-system logical-system-name <target-routing-instance (default | routing-instance-name)>;
  target-routing-instance (default | routing-instance-name);
}
}

```

Hierarchy Level

[edit access]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure a subscriber access profile that includes subscriber access, L2TP, or PPP properties.

Options

profile-name—Name of the profile.

For CHAP, the name serves as the mapping between peer identifiers and CHAP secret keys. This entity is queried for the secret key whenever a CHAP challenge or response is received.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the PPP Authentication Protocol

Configuring Access Profiles for L2TP or PPP Parameters

Configuring L2TP Properties for a Client-Specific Profile

Configuring an L2TP Access Profile on the LNS

Configuring an L2TP LNS with Inline Service Interfaces

Configuring PPP Properties for a Client-Specific Profile

Configuring Service Accounting with JSRC

Configuring Service Accounting in Local Flat Files

AAA Service Framework Overview

[Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management | 428](#)

profile (Captive Portal Content Delivery)

Syntax

```
profile name
  cpcd-rule-sets rule-set-name;
  cpcd-rules rule-name;
  dynamic;
  http-redirect-options url;
  ipda-rewrite-options {
    destination-address destination-address;
    destination-port destination-port;
  }
}
```

Hierarchy Level

[edit services [captive-portal-content-delivery](#)]

Release Information

Statement introduced in Junos OS Release 11.4.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

Description

Configure the service profile for HTTP redirect services, also known as Captive Portal and Content Delivery (CPCD) services. The profile contains rules or rule sets that specify the details of the service. The CPCD service profile is included in a service set that applies the service to a service interface. This statement is supported only for static CPCD; it is not supported for converged services CPCD.

Options

cpcd-rule-sets *rule-set-name*—Specify a list of sets of service rules.

cpcd-rules *rule-name*—Specify a service rule.

dynamic—Indicate that the service is a dynamic, converged service.

http-redirect-options *url*—Redirect the packets to the specified URL. The URL must begin with **http://** or **https://**.

ipda-rewrite-options *destination*—Rewrite the destination for the packets to send them to a remote HTTP redirect server.

- **destination-address *address***—IP address of the remote HTTP redirect server.

- **destination-port *number***—(Optional) Port number for the remote HTTP redirect server. Requires the destination address to be specified as well.

Required Privilege Level

services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

HTTP Redirect Service Overview 473
Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services 484
Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services 494
Configuring Routing Engine-Based, Static HTTP Redirect Services 505
Configuring Routing Engine-Based, Converged HTTP Redirect Services 519

profile (LRF)

Syntax

```

profile profile-name {
  collector collector-name {
    destination {
      address collector-address;
      port collector-port-number;
    }
    source-address source-address;
  }
  http-log-multiple-transactions;
  policy-based-logging;
  rule lrf-rule-name {
    then {
      report {
        collector collector-name;
        template template-name;
        time-limit time-interval;
        volume-limit volume;
      }
    }
  }
  template template-name {
    format ipfix;
    template-tx-interval tx-time;
    template-type template-type;
    trigger-type (session-close | time | volume);
  }
  vendor-support ibm;
}

```

Hierarchy Level

[edit services lrf]

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

vendor-support option introduced in Junos OS Release 17.2.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure an LRF profile to specify a set of logging and reporting parameters, which includes data templates, collectors, and LRF rules.

For Junos OS Subscriber Aware, you can then assign an LRF profile to a subscriber by assigning the profile to the TDF service set associated with the subscriber's TDF domain.

For Junos OS Broadband Subscriber Management, you can then assign the LRF profile to the service set that is configured for application-aware policy control.

Options

profile-name—Name of the LRF profile.

Range: Up to 63 characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers](#) | 460

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management](#) | 459

[Logging and Reporting Function for Subscribers](#) | 441

profile (Services PCEF)

Syntax

```
profile pcef-profile-name;
```

Hierarchy Level

```
[edit services pcef]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Configure a policy and charging enforcement function (PCEF) profile that is a placeholder profile with no configuration options. This profile must be created to enable PCEF functionality on the services plane. You apply this placeholder profile to the subscriber-aware service set.

Options

pcef-profile-name—Name of the PCEF profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Applying Services to Subscriber-Aware Traffic with a Service Set

[Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control](#) | 424

profiles (PCEF)

Syntax

```
profiles profile-name {
  aaa-policy-control {
    aaa-profile aaa-profile-name;
    pcc-rulebases [rulebase-name <time-of-day-profile profile-name>];
    user-password password;
  }
  dynamic-policy-control {
    pcc-rules {
      [rule-name precedence number <time-of-day-profile profile-name>];
    }
    pcc-rulebases {
      [rulebase-name <time-of-day-profile profile-name>];
    }
    diameter-profile gx-profile-name;
  }
  static-policy-control {
    pcc-rules {
      [rule-name precedence number <time-of-day-profile profile-name>];
    }
    pcc-rulebases {
      [rulebase-name <time-of-day-profile profile-name>];
    }
  }
}
```

Hierarchy Level

```
[edit unified-edge pcef],
[edit services pcef]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Set up the overall policy and charging enforcement function (PCEF) configuration that can be applied to subscribers.

NOTE: You can configure only one of the following statements in a PCEF profile:
aaa-policy-control, **static-policy-control**, or **dynamic-policy-control**.

You can configure a maximum of 32 policy and charging control (PCC) rules in a PCEF profile. There is no limit to the number of PCC rulebases you can configure in a PCEF profile.

If you are using Junos OS Subscriber Aware, configure the PCEF profile at the **[edit unified-edge pcef]** hierarchy level. You then assign this profile to the subscriber's TDF domain or to the domain selection configuration.

If you are using Junos OS Broadband Subscriber Management, configure the PCEF profile at the **[edit services pcef]** hierarchy level. The **static-policy-control** option is applicable to PCC rule activation through a dynamic profile, and you assign the PCEF profile to the dynamic profile. Starting in Junos OS Release 18.2R1, the **dynamic-policy-control** option is also available and is applicable to direct rule activation by a policy and charging rules function (PCRF) server; you assign the PCEF profile to the access profile. The **aaa-policy-control** option is not applicable to subscriber management.

Options

profile-name—Name of the PCEF profile.

Range: 1 through 63 characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Dynamic Policies

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Static Policies

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Policies That a RADIUS Server Controls

[Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management](#) | 423

profile-type (Dynamic Service Profiles)

Syntax

```
profile-type remote-device-service;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Specify that the dynamic service profile containing this statement is not applied locally to the router. Instead, it is applied to an external device by means of the remote device services manager daemon (rdmd). It enables an external authority, such as PCRF to reference the dynamic service profile to provision or deprovision services (charging rules) on the remote device.

The content of the service is limited to a set of variables that are translated to NETCONF XML protocol remote procedure calls (RPCs) by a dictionary provisioned on the MX series router. The dictionary contains entries corresponding to different service profiles. Each entry contains the RPCs that correspond to the service variables and install the service on the remote device.

The **remote-device-service** type effectively separates the external authority from how and where a service is provisioned. The external authority simply references the profile name; if the service profile has this type, then the service is provisioned on a remote device. If the profile does not have this type, then the service is provisioned on the router.

Options

remote-device-service—Service profile type that causes the profile to be applied to an external device by means of the remote device services manager daemon (rdmd).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Remote Device Management for Service Provisioning](#) | 648

[Remote Device Services Manager \(RDSM\) Overview](#) | 631

promiscuous-mode (Dynamic IGMP Interface)

Syntax

```
promiscuous-mode;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name],
```

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Specify that the interface accepts IGMP reports from hosts on any subnetwork. When you enable promiscuous mode, all routing devices on the Ethernet segment must be configured for promiscuous mode. 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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Accepting IGMP Messages from Remote Subnetworks

protocol (Application Identification)

Syntax

```
protocol (http | ssl | tcp | udp);
```

Hierarchy Level

```
[edit services application-identification application application-name over protocol-type signature I4-I7-signature-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Identify the protocol bundles to be monitored to classify applications. This statement is not available if the MX Series router is running Next Gen Services.

Options

http—Use the HTTP protocol .

ssl—Use the SSL protocol.

tcp—Use the TCP protocol.

udp—Use the UDP protocol.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

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](#) | 40

[Configuring Schedulers in a Dynamic Profile for Subscriber Access](#) | 52

protocol (Flow Descriptions)

Syntax

```
protocol number;
```

Hierarchy Level

```
[edit unified-edge pcef flow-description flow-identifier],  
[edit services pcef flow-description flow-identifier]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef flow-description *flow-identifier*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Specify a protocol type to identify the subscriber traffic that you want the service data flow (SDF) filter to detect. If you specify the **protocol** statement, you must specify a protocol number.

If you are using Junos OS Subscriber Aware, specify the protocol type at the **[edit unified-edge pcef flow-description *flow-identifier*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the protocol type at the **[edit services pcef flow-description *flow-identifier*]** hierarchy level.

Default

If you do not specify the **protocol** statement, the default is any protocol.

Options

number—Number that specifies the IP protocol type.

Range: 1 through 255

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Service Data Flow Filters | 412](#)

[Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment](#)

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

protocols (DDoS)

Syntax

```

protocols protocol-group (aggregate | packet-type) {
    bandwidth packets-per-second;
    burst size;
    bypass-aggregate;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection-mode (automatic | off | on);
    flow-detect-time seconds;
    flow-level-bandwidth {
        logical-interface flow-bandwidth;
        physical-interface flow-bandwidth;
        subscriber flow-bandwidth;
    }
    flow-level-control {
        logical-interface flow-control-mode;
        physical-interface flow-control-mode;
        subscriber flow-control-mode;
    }
    flow-level-detection {
        logical-interface flow-operation-mode;
        physical-interface flow-operation-mode;
        subscriber flow-operation-mode;
    }
    flow-recover-time seconds;
    flow-timeout-time seconds;
    fpc slot-number {
        bandwidth-scale percentage;
        burst-scale percentage;
        disable-fpc;
    }
    no-flow-logging
    priority level;
    recover-time seconds;
    timeout-active-flows;
}

```

Hierarchy Level

[edit system [ddos-protection](#)]

Release Information

Statement introduced in Junos OS Release 11.2.

Statement introduced in Junos OS Release 12.3R2 on EX9200 switches and T4000 routers.

Support for Enhanced Subscriber Management added in Junos OS Release 17.3R1.

Description

(MX Series routers with only MPCs, T4000 Core Routers with only FPC5s, or EX9200 switches) Configure control plane DDoS protection policers for all supported packet types within a protocol group or for a particular supported packet type within a protocol group.

NOTE: For the available control plane DDoS protection policer configuration options on PTX Series routers and QFX Series switches, which are different from the options described here, see *protocols (DDoS) (PTX Series and QFX Series)*.

Options

aggregate—Configure the policer to monitor all control packets within the protocol group. You can configure an aggregate policer for any protocol group.

packet-type—(Optional) Name of the control packet type to be policed. You can configure a specific policer for only the following packet types and protocol groups:

- **arp**—The following ARP packet types are available:
 - **aggregate**—Applies to the combination of all types of control packet traffic for this protocol group.
- **bgp**—The following BGP packet types are available:
 - **aggregate**—Applies to the combination of all types of control packet traffic for this protocol group.
- **bgpv6**—The following BGPv6 packet types are available:
 - **aggregate**—Applies to the combination of all types of control packet traffic for this protocol group.
- **dhcpv4**—The following packet types are available for DHCPv4 traffic:
 - **ack**—DHCPACK packets.
 - **bad-packets**—DHCPv4 packets with bad formats.
 - **bootp**—DHCPBOOTP packets.
 - **decline**—DHCPDECLINE packets.
 - **discover**—DHCPDISCOVER packets.
 - **force-renew**—DHCPFORCERENEW packets.
 - **inform**—DHCPINFORM packets.
 - **lease-active**—DHCPLEASEACTIVE packets.
 - **lease-query**—DHCPLEASEQUERY packets.
 - **lease-unassigned**—DHCPLEASEUNASSIGNED packets.
 - **lease-unknown**—DHCPLEASEUNKNOWN packets.
 - **nak**—DHCPNAK packets.
 - **no-message-type**—DHCP packets that are missing the message type.
 - **offer**—DHCP OFFER packets.
 - **release**—DHCPRELEASE packets.
 - **renew**—DHCPRENEW packets.
 - **request**—DHCPREQUEST packets.
 - **unclassified**—All unclassified packets in the protocol group.
- **dhcpv6**—The following packet types are available for DHCPv6 traffic:

- **advertise**—ADVERTISE packets.
- **confirm**—CONFIRM packets.
- **decline**—DECLINE packets.
- **information-request**—INFORMATION-REQUEST packets.
- **leasequery**—LEASEQUERY packets.
- **leasequery-data**—LEASEQUERY-DATA packets.
- **leasequery-done**—LEASEQUERY-DONE packets.
- **leasequery-reply**—LEASEQUERY-REPLY packets.
- **rebind**—REBIND packets.
- **reconfigure**—RECONFIGURE packets.
- **relay-forward**—RELAY-FORWARD packets.
- **relay-reply**—RELAY-REPLY packets.
- **release**—RELEASE packets.
- **renew**—RENEW packets.
- **reply**—REPLY packets.
- **request**—REQUEST packets.
- **solicit**—SOLICIT packets.
- **unclassified**—All unclassified packets in the protocol group.
- **filter-action**—The following packet types are available for unclassified firewall filter action packets, sent to the host because of reject terms in firewall filters:
 - **filter-v4**—Unclassified IPv4 filter action packets.
 - **filter-v6**—Unclassified IPv6 filter action packets.
 - **other**—All other unclassified filter action packets that are not IPv4 or IPv6.
- **frame-relay**—The following packet types are available for Frame Relay traffic:
 - **frf15**—Multilink frame relay FRF.15 packets.
 - **frf16**—Multilink frame relay FRF.16 packets.
- **ip-fragments**—The following packet types are available for IP fragments:
 - **first-fragment**—First IP fragment.
 - **trail-fragment**—Last IP fragment.

- **ip-options**—The following packet types are available for IP option traffic:
 - **non-v4v6**—Options packets other than IPv4/v6.
 - **router-alert**—Router alert options packets.
 - **unclassified**—All unclassified packets in the protocol group.
- **l2tp**—The following packet types are available for L2TP traffic:
 - **cdn**—Call-Disconnect-Notify message packets.
 - **hello**—Hello message packets.
 - **iccn**—Incoming-Call-Connected message packets.
 - **icrq**—Incoming-Call-Request message packets.
 - **scccn**—Start-Control-Connection-Connected message packets.
 - **sccrq**—Start-Control-Connection-Request message packets.
 - **stopccn**—Stop-Control-Connection-Notification message packets.
 - **unclassified**—All unclassified packets in the protocol group.
- **mcast-snoop**—Control traffic for multicast snooping.
 - **igmp**—Snooped IGMP traffic.
 - **mld**—Snooped MLD traffic.
 - **pim**—Snooped PIM control traffic.
- **mlp**—The following MLP packet types are available:
 - **aggregate**—Applies to the combination of all types of control packet traffic for this protocol group.
 - **add**—Add requests; internal MAC address learning request packets sent to the host.
 - **delete**—Delete requests; internal MAC address learning request packets sent to the host.
 - **lookup**—Lookup requests; internal MAC address learning request packets sent to the host.
 - **unclassified**—All unclassified packets in the protocol group.
 - **macpin-exception**—Exceptions to MAC address pinning (wherein dynamically learned MAC addresses are pinned to prevent looping caused by MAC moves from duplicate MAC detection).

- **ndpv6**—The following NDPv6 packet types are available, except where noted, starting in 14.1R8, 14.2R8, 15.1R5, 15.1F7, and 16.1R1:
 - **aggregate**—Applies to the combination of all types of control packet traffic for this protocol group.
 - **invalid-hop-limit**—(Starting in 16.1R2) Invalid hop limit packets. These messages might represent crafted packets in a malicious network-based packet flood.
 - **neighbor-advertisement**—Neighbor advertisement packets. These are messages used for duplicate address detection and to test reachability of neighbors. Neighbor advertisements are sent in response to neighbor solicitation messages.
 - **neighbor-solicitation**—Neighbor solicitation packets. These are messages used for duplicate address detection and to test reachability of neighbors.
 - **redirect**—Redirect packets.
 - **router-advertisement**—Router advertisement packets. These are messages sent to announce the presence of the router, advertise prefixes, assist in address configuration, and share other link information such as MTU size and hop limit. The IPv6 nodes on the link can use this information to configure themselves with an IPv6 address and routing information such as the default gateway.
 - **router-solicitation**—Router solicitation packets. These are messages sent by IPv6 nodes when they come online to solicit immediate router advertisements from the router.
- **ppp**—The following PPP packet types are available:
 - **authentication**—PPP authentication protocol packets.
 - **echo-rep**—LCP echo reply packets.
 - **echo-req**—LCP echo request packets.
 - **ipcp**—IP Control Protocol packets.
 - **ipv6cp**—IPv6 Control Protocol packets.
 - **isis**—IS-IS packets.
 - **lcp**—Link Control Protocol packets.
 - **mlppp-lcp**—MLPPP LCP packets.
 - **mplscp**—MPLS Control Protocol packets.
 - **unclassified**—All unclassified packets in the protocol group.

- **pppoe**—The following PPPoE packet types are available:
 - **padi**—PADI packets.
 - **padm**—PADM packets.
 - **padn**—PADN packets.
 - **pado**—PADO packets.
 - **padr**—PADR packets.
 - **pads**—PADS packets.
 - **padt**—PADT packets.
- **radius**—The following RADIUS packet types are available:
 - **accounting**—RADIUS accounting packets.
 - **authorization**—RADIUS authorization packets.
 - **server**—RADIUS server traffic.
 - **unclassified**—All unclassified packets in the protocol group.
- **re-services**—The following packet type is available for Routing Engine-based HTTP redirect IPv4 traffic:
 - **captive-portal**—Routing Engine-based captive portal content delivery packets.
- **re-services-v6**—The following packet type is available for Routing Engine-based HTTP redirect IPv6 traffic:
 - **captive-portal**—Routing Engine-based captive portal content delivery packets.
- **resolve**—The following packet types are available for unclassified resolve packets, which are sent to the host because of a traffic request resolve action:
 - **mcast-v4**—Unclassified IPv4 multicast resolve packets.
 - **mcast-v6**—Unclassified IPv6 multicast resolve packets.
 - **ucast-v4**—Unclassified IPv4 unicast resolve packets.
 - **ucast-v6**—Unclassified IPv6 unicast resolve packets.
 - **other**—All other unclassified resolve packets.

- **sample**—The following sample packet types are available:
 - **host**—Host packets.
 - **pfe**—Packet Forwarding Engine packets.
 - **syslog**—System log message packets.
 - **tap**—TAP packets.
- **tcp-flags**—The following TCP-flagged packet types are available:
 - **established**—TCP packets with ACK or RST flags set.
 - **initial**—TCP packets with SYN flag set and ACK flag not set.
 - **unclassified**—TCP packets with flags set any other way than the established and initial packets.
- **unclassified**—The following unclassified packet types are available:
 - **control-layer2**—Unclassified layer 2 control packets.
 - **control-v4**—Unclassified IPv4 control packets.
 - **control-v6**—Unclassified IPv6 control packets.
 - **fw-host**—Unclassified send-to-host firewall packets.
 - **host-route-v4**—Unclassified IPv4 routing protocol and host packets in traffic sent to the router local interface address.
 - **host-route-v6**—Unclassified IPv6 routing protocol and host packets in traffic sent to the router local interface address.
 - **other**—All unclassified packets that do not belong to another type.
- **virtual-chassis**—The following packet types are available for virtual chassis packets:
 - **control-low**—Low-priority control packets.
 - **control-high**—High-priority control packets.
 - **unclassified**—All unclassified packets in the protocol group.
 - **vc-packets**—All exception packets on the virtual chassis link.
 - **vc-ttl-errors**—Virtual chassis TTL error packets.

protocol-group—Name of the protocol group for which traffic is policed. You can configure a policer for any of the following protocol groups:

- **amtv4**—IPv4 AMT traffic.
- **amtv6**—IPv6 AMT traffic.
- **ancp**—ANCP traffic.
- **ancpv6**—ANCPv6 traffic.
- **arp**—ARP traffic.
- **atm**—ATM traffic.
- **bfd**—BFD traffic.
- **bfdv6**—BFDv6 traffic.
- **bgp**—BGP traffic.
- **bgpv6**—BGPv6 traffic.
- **control**—Control traffic.
- **demux-autosense**—Demux autosensing traffic.
- **dhcpv4**—DHCPv4 traffic.
- **dhcpv6**—DHCPv6 traffic.
- **diameter**—Diameter and Gx-Plus traffic.
- **dns**—DNS traffic.
- **dtcp**—DTCP traffic.
- **dynamic-vlan**—Dynamic VLAN exception traffic.
- **egpv6**—EGPv6 traffic.
- **eoam**—EOAM traffic.
- **esmc**—ESMC traffic.
- **fab-probe**—Fab out probe packets.
- **filter-action**—IPv4 and IPv6 firewall filter action packets sent to the host because of reject terms in firewall filters
- **frame-relay**—Frame relay traffic.
- **ftp**—FTP traffic.
- **ftpv6**—FTPv6 traffic.
- **gre**—GRE traffic.
- **icmp**—ICMP traffic.

- **igmp**—IGMP traffic
- **igmpv4v6**—IGMP v4/v6 traffic.
- **igmpv6**—IGMPv6 traffic.
- **inline-ka**—Inline service interfaces keepalive traffic.
- **inline-svcs**—Inline services traffic.
- **ip-fragments**—IP fragments traffic.
- **ip-options**—IP traffic with IP packet header options.
- **isis**—IS-IS traffic.
- **jfm**—JFM traffic.
- **l2pt**—Layer 2 protocol tunneling traffic.
- **lACP**—LACP traffic.
- **ldp**—LDP traffic.
- **ldpv6**—LDPv6 traffic.
- **lldp**—LLDP traffic.
- **lmp**—LMP traffic.
- **lmpv6**—LMPv6 traffic.
- **mac-host**—Layer 2 MAC send-to-host traffic.
- **mcast-snoop**—Control traffic for multicast snooping.
- **mlp**—MLP traffic.
- **msdp**—MSDP traffic.
- **msdpv6**—MSDPv6 traffic.
- **multicast-copy**—Host copy traffic due to multicast routing.
- **mvrp**—MVRP traffic.
- **ndpv6**—NDPv6 traffic.
- **ntp**—NTP traffic.
- **oam-lfm**—OAM-LFM traffic.
- **ospf**—OSPF traffic.
- **ospfv3v6**—OSPFv3/IPv6 traffic.
- **pfe-alive**—Packet Forwarding Engine keepalive traffic.
- **pim**—PIM traffic.

- **pimv6**—PIMv6 traffic.
- **pmvrp**—PMVRP traffic.
- **pos**—POS traffic.
- **ppp**—PPP traffic.
- **pppoe**—PPPoE traffic.
- **ptp**—PTP traffic.
- **pvstp**—PVSTP traffic.
- **radius**—RADIUS traffic.
- **re-services**—Captive portal content delivery IPv4 traffic for Routing Engine HTTP redirect.
- **re-services-v6**—Captive portal content delivery IPv6 traffic for Routing Engine HTTP redirect.
- **redirect**—Traffic that triggers ICMP redirects.
- **reject**—Packets rejected by a next-hop forwarding decision.
- **rejectv6**—V6 packets rejected by a next-hop forwarding decision.
- **resolve**—Unclassified IPv4 and IPv6 resolve packets sent to the host because of a traffic request resolve action.
- **rip**—RIP traffic.
- **ripv6**—RIPv6 traffic.
- **rsvp**—RSVP traffic.
- **rsvpv6**—RSVPv6 traffic.
- **services**—Service traffic.
- **snmp**—SNMP traffic.
- **snmpv6**—SNMPv6 traffic.
- **ssh**—SSH traffic.
- **sshv6**—SSHv6 traffic.
- **stp**—STP traffic.
- **syslog**—System log messages UDP traffic on port 6333 for the Routing Engine syslog server.
- **tacacs**—TACACS traffic.
- **tcp-flags**—Traffic with TCP flags.
- **telnet**—TELNET traffic.
- **telnetv6**—TELNETv6 traffic.
- **ttl**—TTL traffic.

- **tunnel-fragment**—Tunnel fragments traffic.
- **tunnel-ka**—Tunnel keepalive traffic.
- **unclassified**—Unclassified traffic.
- **virtual-chassis**—Virtual chassis traffic.
- **vrrp**—VRRP traffic.
- **vrrpv6**—VRRPv6 traffic.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers

Understanding Distributed Denial-of-Service Protection with IPv6 Neighbor Discovery Protocol

protocols (Dynamic Profiles)

Syntax

```

protocols {
  igmp {
    interface interface-name {
      accounting;
      disable;
      group-limit limit;
      group-policy;
      group-threshold value;
      immediate-leave
      log-interval seconds;
      no-accounting;
      oif-map;
      passive;
      promiscuous-mode;
      ssm-map ssm-map-name;
      ssm-map-policy ssm-map-policy-name
      static {
        group group {
          source source;
        }
      }
      version version;
    }
  }
  mld {
    interface interface-name {
      (accounting | no-accounting);
      disable;
      group-limit limit;
      group-policy;
      group-threshold value;
      immediate-leave;
      log-interval seconds;
      oif-map;
      passive;
      ssm-map ssm-map-name;
      ssm-map-policy ssm-map-policy-name;
      static {
        group multicast-group-address {
          exclude;
          group-count number;
        }
      }
    }
  }
}

```

```

    group-increment increment;
    source ip-address {
        source-count number;
        source-increment increment;
    }
}
}
}
version version;
}
}
router-advertisement {
    interface interface-name {
        current-hop-limit number;
        default-lifetime seconds;
        (managed-configuration | no-managed-configuration);
        max-advertisement-interval seconds;
        min-advertisement-interval seconds;
        (other-stateful-configuration | no-other-stateful-configuration);
        prefix prefix;
        reachable-time milliseconds;
        retransmit-timer milliseconds;
    }
}
}
}

```

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 Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) Protocol on the router and configure interface-specific values on dynamic interfaces for each protocol.

PIM and MLD manage multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use these protocols to learn which groups have members on each of their attached physical networks. Enable IGMP for the router to receive IPv4 or IPv6 multicast traffic. Enable MLD for the router to receive IPv6 multicast traffic. MLD is needed only for IPv6 networks.

You can also use this statement to enable router advertisement for IPv6 Neighbor Discovery protocol and configure interface-specific values on dynamic interfaces.

You can configure these protocols in a dynamic client profile or a dynamic service profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Default

IGMP is disabled on the router. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring IGMP

Configuring MLD

provisioning-method (Remote Device Management)

Syntax

```
provisioning-method {
  netconf {
    bulk-interval milliseconds;
    bulk-limit number;
    connection-retry-interval seconds;
    password password;
    port port-number;
    reconfigure-bulk-limit number;
    response-timeout seconds;
    response-timeout-count number;
    user-name name;
  }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name system services remote-device-management service-device
device-name],
[edit system services remote-device-management service-device device-name]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Configure the method for provisioning and deprovisioning services on the remote device. The NETCONF XML management protocol is currently the only supported method.

The remaining statement is explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

RELATED DOCUMENTATION

[Configuring Remote Device Management for Service Provisioning](#) | 648

[Remote Device Services Manager \(RDSM\) Overview](#) | 631

radius (Access Profile)

Syntax

```
radius {
  accounting-server [ ip-address ];
  attributes {
    exclude
      attribute-name packet-type;
    standard-attribute number {
      packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
    }
    vendor-id id-number {
      vendor-attribute vsa-number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
      }
    }
  }
  ignore {
    dynamic-iflset-name;
    framed-ip-netmask;
    idle-timeout;
    input-filter;
    logical-system-routing-instance;
    output-filter;
    session-timeout;
    standard-attribute number;
    vendor-id id-number {
      vendor-attribute vsa-number;
    }
  }
}
authentication-server [ ip-address ];
options {
  accounting-session-id-format (decimal | description);
  calling-station-id-delimiter delimiter-character;
  calling-station-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    nas-identifier;
  }
  chap-challenge-in-request-authenticator;
  client-accounting-algorithm (direct | round-robin);
  client-authentication-algorithm (direct | round-robin);
```

```

coa-dynamic-variable-validation;
ethernet-port-type-virtual;
interface-description-format {
    exclude-adapter;
    exclude-channel;
    exclude-sub-interface;
}
ip-address-change-notify message;
juniper-access-line-attributes;
nas-identifier identifier-value;
nas-port-extended-format {
    adapter-width width;
    ae-width width;
    port-width width;
    slot-width width;
    stacked-vlan-width width;
    vlan-width width;
    atm {
        adapter-width width;
        port-width width;
        slot-width width;
        vci-width width;
        vpi-width width;
    }
}
nas-port-id-delimiter delimiter-character;
nas-port-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    nas-identifier;
    order {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        interface-text-description;
        nas-identifier;
        postpend-vlan-tags;
    }
    postpend-vlan-tags;
}

```

```

nas-port-type {
    ethernet {
        port-type;
    }
}
override {
    calling-station-id remote-circuit-id;
    nas-ip-address tunnel-client-gateway-address;
    nas-port tunnel-client-nas-port;
    nas-port-type tunnel-client-nas-port-type;
}
remote-circuit-id-delimiter;
remote-circuit-id-fallback;
remote-circuit-id-format {
    agent-circuit-id;
    agent-remote-id;
}
revert-interval interval;
service-activation {
    dynamic-profile (optional-at-login | required-at-login);
    extensible-service (optional-at-login | required-at-login);
}
vlan-nas-port-stacked-format;
}
preauthentication-server ip-address;
}

```

Hierarchy Level

[edit access [profile](#) *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description

Configure the RADIUS parameters that the router uses for AAA authentication and accounting for subscribers.

Options

accounting-server—(MX Series only) Specify a list of the RADIUS accounting servers used for accounting for DHCP, L2TP, and PPP clients.

Values: *ip-address*—IP version 4 (IPv4) address.

authentication-server—(SRX Series only) Specify a list of the RADIUS authentication servers used to authenticate DHCP, L2TP, and PPP clients. The servers in the list are also used as RADIUS dynamic-request servers, from which the router accepts and processes RADIUS disconnect requests, CoA requests, and dynamic service activations and deactivations.

Values: *ip-address*—IPv4 address.

preauthentication-server—(MX Series only) Starting in Junos OS Release 13.3, specify the RADIUS preauthentication server, which is used for the LLID service.

NOTE: You cannot configure this statement if the Calling-Station-ID attribute is excluded from RADIUS Access-Request messages by the **exclude** statement.

Values: *ip-address*—IPv4 address.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Authentication and Accounting Parameters for Subscriber Access](#)

[RADIUS Logical Line Identifier \(LLID\) Overview](#)

[Configuring Logical Line Identification \(LLID\) Preauthentication](#)

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	183
Configuring CoS Adjustment Control Profiles	186
Verifying the CoS Adjustment Control Profile Configuration	188
adjustment-control-profiles	712
application (Adjustment Control Profiles)	729

radius-flow-tap

Syntax

```
radius-flow-tap {
    forwarding-class class-name;
    interfaces interface-name;
    logical-system name <routing-instance routing-instance>;
    multicast-interception;
    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;
                }
            }
        }
    }
    snmp {
        notify-targets ip-address;
    }
    routing-instance routing-instance-name;
    source-ipv4-address ipv4-address;
}
```

Hierarchy Level

```
[edit services]
```

Release Information

Statement introduced in Junos OS Release 9.4.

Description

Configure the radius-flow-tap service for subscriber secure policy mirroring. Both RADIUS-initiated and Dynamic Tasking Control Protocol (DTCP)-initiated mirroring are supported.

Starting in Junos OS Release 17.3R1, the radius-flow-tap service can run concurrently on the same router with the FlowTapLite service. The FlowTapLite service is a version of the flow-tap service ([[edit services flow-tap](#)]) that is configured only on tunnel interfaces on MX Series routers.

In earlier releases, the radius-flow-tap and FlowTapLite services cannot run concurrently on an MX Series router, preventing you from running FlowTapLite monitoring and subscriber secure policy mirroring at the same time.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

- flow-tap—To view this statement in the configuration.
- flow-tap-control—To add this statement to the configuration.

Release History Table

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, the radius-flow-tap service can run concurrently on the same router with the FlowTapLite service.

RELATED DOCUMENTATION

Configuring Support for Subscriber Secure Policy Mirroring 574
Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring 597
Subscriber Secure Policy Overview 553

radius-server

Syntax

```
radius-server server-address {
    accounting-port port-number;
    accounting-retry number;
    accounting-timeout seconds;
    dynamic-request-port port-number;
    max-outstanding-requests value;
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level

```
[edit access],
[edit access profile profile-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

max-outstanding-requests introduced in Junos OS Release 11.4.

accounting-retry and **accounting-timeout** introduced in Junos OS Release 14.1.

dynamic-request-port option added in Junos OS Release 14.2R1 for MX Series routers.

preauthentication-port and **preauthentication-secret** options added in Junos OS Release 15.1 for MX Series routers.

accounting-port introduced in Junos OS Release 13.2X50-D10 for EX Series switches with support for Enhanced Layer 2 software (ELS). It was introduced in Junos OS without ELS in the following releases: Junos OS Releases 12.3R10, 14.1X53-D25, and 15.1R4 for EX Series switches.

Support for IPv6 **server-address** introduced in Junos OS Release 16.1.

Description

Configure RADIUS for subscriber access management, L2TP, or PPP.

To configure multiple RADIUS servers, include multiple **radius-server** statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options

server-address—IPv4 or IPv6 address of the RADIUS server.

accounting-port—(EX Series, M Series, MX Series, PTX Series, T Series only) Configure the port number on which to contact the RADIUS accounting server. This statement was introduced in Junos OS Release 13.2X50-D10 for EX Series switches with support for Enhanced Layer 2 software (ELS). It was introduced in Junos OS without ELS in the following releases: Junos OS Releases 12.3R10, 14.1X53-D25, and 15.1R4 for EX Series switches.

NOTE: Specifying the accounting port is optional, and port 1813 is the default. However, we recommend that you configure it in order to avoid confusion, as some RADIUS servers might refer to an older default.

Values: *port-number*—Port number on which to contact the RADIUS accounting server. Most RADIUS servers use port 1813, as specified in RFC 2866.

Default: 1813

accounting-retry—(MX Series, T Series only) Starting in Junos OS Release 14.1, configure the number of times the device retransmits RADIUS accounting messages when no response is received from the server. When you do not configure this statement, the number of retry attempts is determined by the **retry** statement.

NOTE: To successfully set a retry limit for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *number*—Number of retry attempts.

Range: 0 through 100

Default: 0 (disabled)

accounting-timeout—(MX Series, T Series only) Starting in Junos OS Release 14.1, configure how long the local device waits to receive a response from a RADIUS accounting server before retransmitting the message. When you do not configure this statement, the length of the timeout is determined by the **timeout** statement.

NOTE: To successfully set a timeout value for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *seconds*—Duration of timeout period.

Range: 0 through 1000 seconds

Default: 0 (disabled)

dynamic-request-port—(MX Series only) Starting in Junos OS Release 14.2R1, specify the port that the router monitors for dynamic (CoA) requests from the specified RADIUS servers. You can configure a port globally or for a specific access profile.

You must either use the default port for all RADIUS servers or configure the same nondefault port for all RADIUS servers. This rule applies at both the global access and access profile levels.

NOTE: Any other configuration results in a commit check failure. Multiple port numbers—that is, different port numbers for different servers—are not supported.

Values: *port-number*—Number of the monitored port.

Default: 3799 (as specified in RFC 5176)

max-outstanding-requests—(MX Series only) Starting in Junos OS Release 11.4, configure the maximum number of outstanding requests for this RADIUS server. An increase in this value is immediate while a decrease is more gradual if the current number of outstanding requests exceeds the new value.

Values: *requests*—Maximum number of outstanding requests for this RADIUS server.

Range: 0 through 2000 outstanding requests per server

Default: 1000 outstanding requests per server

port—(EX Series, M Series, MX Series, SRX Series, T Series only) Configure the port number on which to contact the RADIUS server.

Values: *port-number*—Port number on which to contact the RADIUS server.

Default: 1812 (as specified in RFC 2865)

preauthentication-port—(MX Series only) Starting in Junos OS Release 15.1 for MX Series routers, configure the port number on which to contact the RADIUS server for logical line identification (LLID) preauthentication requests. If you do not configure a separate UDP port for preauthentication purposes, the same UDP port that you configure for authentication messages by including the **port port-number** statement is used.

Values: *port-number*—Port number used for preauthentication requests to contact the RADIUS server.

preauthentication-secret—(MX Series only) Starting in Junos OS Release 15.1 for MX Series routers, configure the password to use with the RADIUS server for LLID preauthentication requests. If you do not configure a separate UDP password for preauthentication purposes, the same password that you configure for authentication messages by including the **secret password** statement is used. The secret password used by the local router must match that used by the server.

Values: *password*—Password to use. To include spaces enclose the character string in quotation marks.

retry—(EX Series, M Series, MX Series, PTX Series, T Series only) Specify the number of times that the device is allowed to attempt to contact a RADIUS authentication or accounting server. You can override the retry limit for accounting servers with the **accounting-retry** statement.

NOTE: To successfully set a retry limit for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *attempts*—Number of times that the router is allowed to attempt to contact a RADIUS server.

Range: 1 through 100

Default: 3

routing-instance—(SRX Series, vSRX only) Configure the routing instance used to send RADIUS packets to the RADIUS server.

Values: *routing-instance-name*—Routing instance name.

source-address—(SRX Series, vSRX only) Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. Support for IPv6 **source-address** was introduced in Junos OS Release 16.1.

Values: *source-address*—Valid IPv4 or IPv6 address configured on one of the router or switch interfaces. On M Series routers only, the source address can be an IPv6 address and the UDP source port is 514.

timeout—(SRX Series, vSRX only) Configure the amount of time that the local device waits to receive a response from RADIUS authentication and accounting servers. You can override the timeout value for accounting servers with the **accounting-timeout** statement.

NOTE: To successfully set a timeout value for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *seconds*—Amount of time to wait.

Range: 1 through 1000 seconds

Default: 3 seconds

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Authentication and Accounting Parameters for Subscriber Access](#)

[Configuring the PPP Authentication Protocol](#)

[Configuring RADIUS Authentication for L2TP](#)

[Configuring RADIUS System Accounting](#)

[Configuring RADIUS-Initiated Dynamic Request Support | 5](#)

[RADIUS Logical Line Identifier \(LLID\) Overview](#)

[RADIUS Attributes for LLID Preauthentication Requests](#)

[show network-access aaa statistics](#)

[clear network-access aaa statistics](#)

rate-limit

Syntax

```
rate-limit limit;
```

Hierarchy Level

```
[edit system services finger],  
[edit system services ftp],  
[edit system services netconf ssh],  
[edit system services ssh],  
[edit system services telnet],  
[edit system services tftp-server],  
[edit system services xnm-clear-text],  
[edit system services xnm-ssl]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Configure the maximum number of connections attempts per minute, per protocol (either IPv6 or IPv4) on an access service. For example, a rate limit of 10 allows 10 IPv6 telnet session connection attempts per minute and 10 IPv4 telnet session connection attempts per minute.

Default

150 connections

Options

rate-limit *limit*—(Optional) Maximum number of connection attempts allowed per minute, per IP protocol (either IPv4 or IPv6).

Range: 1 through 250

Default: 150

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring clear-text or SSL Service for Junos XML Protocol Client Applications](#)

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](#)

redirect (PCC Action Profiles)

Syntax

```
redirect {  
    url url-name;  
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],  
[edit services pcef pcc-action-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify HTTP redirection to a URL. If you configure this, the PCC action profile can only be used in PCC rules that match only HTTP-based applications and all flows.

If you are using Junos OS Subscriber Aware, specify the redirection at the **[edit unified-edge pcef pcc-action-profiles *profile-name*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the redirection at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

[Configuring Policy and Charging Control Action Profiles for Subscriber Management](#) | 416

remote-address

Syntax

```
remote-address (ipv4-address ipv4-address | ipv6-address ipv6-address);
```

Hierarchy Level

```
[edit unified-edge pcef flow-descriptions flow-identifier],  
[edit services pcef flow-descriptions flow-identifier]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef flow-descriptions *flow-identifier*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Specify a remote IP address for the service data flow (SDF) filter.

If you are using Junos OS Subscriber Aware, specify the remote IP address at the **[edit unified-edge pcef flow-descriptions *flow-identifier*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the remote IP address at the **[edit services pcef flow-descriptions *flow-identifier*]** hierarchy level.

Options

ipv4-address—IPv4 address.

ipv6-address—IPv6 address.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Service Data Flow Filters](#) | 412

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

remote-device-management

Syntax

```
remote-device-management {
  service-device device-name {
    access-domain {
      vlan-id-list [vlan-id-low-vlan-id-high vlan-id]
    }
    address ip-address;
    dictionary absolute file path;
    provisioning-method {
      netconf {
        bulk-interval milliseconds;
        bulk-limit number;
        connection-retry-interval seconds;
        password password;
        port port-number;
        reconfigure-bulk-limit number;
        response-timeout seconds;
        response-timeout-count number;
        user-name name;
      }
    }
  }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name system services],
[edit system services]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

RELATED DOCUMENTATION

[Configuring Remote Device Management for Service Provisioning](#) | 648

[Remote Device Services Manager \(RDSM\) Overview](#) | 631

remote-port-range

Syntax

```
remote-port-range {  
    low low-value;  
    high high-value;  
}
```

Hierarchy Level

```
[edit unified-edge pcef flow-descriptions flow-identifier],  
[edit services pcef flow-descriptions flow-identifier]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef flow-descriptions *flow-identifier*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Specify the remote port range to identify the subscriber traffic that you want the service data flow (SDF) filter to detect.

If you are using Junos OS Subscriber Aware, specify the remote port range at the **[edit unified-edge pcef flow-descriptions *flow-identifier*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the remote port range at the **[edit services pcef flow-descriptions *flow-identifier*]** hierarchy level.

NOTE: You can specify either a remote port range or a list of remote ports, but not both.

Default

If you configure neither the **remote-port-range** nor the **remote-ports** statement, the default is any remote port.

Options

high-value—Upper boundary for the remote port range.

Range: 1 through 65,535

low-value—Lower boundary for the remote port range.

Range: 1 through 65,535

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Service Data Flow Filters | 412](#)

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

remote-ports

Syntax

```
remote-ports [number];
```

Hierarchy Level

```
[edit unified-edge pcef flow-description flow-identifier],  
[edit services pcef flow-description flow-identifier]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef flow-descriptions *flow-identifier*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Specify a remote port or list of remote ports to identify the subscriber traffic that you want the service data flow (SDF) filter to detect.

If you are using Junos OS Subscriber Aware, specify the remote ports at the **[edit unified-edge pcef flow-descriptions *flow-identifier*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the remote ports at the **[edit services pcef flow-descriptions *flow-identifier*]** hierarchy level.

NOTE: You can specify either a list of remote ports or a remote port range, but not both.

Default

If you configure neither the **remote-ports** nor the **remote-port-range** statement, the default is any remote port.

Options

number—Port number or list of port numbers. You can specify a maximum of three port numbers in a list.

Range: 1 through 65,535

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Service Data Flow Filters | 412](#)

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

report (LRF Rule)

Syntax

```
report {  
  collector collector-name;  
  template template-name;  
  time-limit time-interval;  
  volume-limit volume;  
}
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name then]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure the actions to take if the LRF rule is matched.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers](#) | 460

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management](#) | 459

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access](#) | 40

rewrite-rules

routing-engine-services

Syntax

```
routing-engine-services;
```

Hierarchy Level

```
[edit services service-set service-set service-set-options]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

Description

When configuring a Routing Engine-based captive portal service, specify the service set options to apply to a service set. The services interfaces on the Routing Engine are identified with an si- prefix (for example, si-1/1/0). The si- interface contains all redirect and rewrite traffic and services for the Routing Engine.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview | 473](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

routing-options (Dynamic Profiles)

Syntax

```

routing-options {
  access {
    route prefix {
      metric route-cost;
      next-hop next-hop;
      preference route-distance;
      tag route-tag;
      tag2 route-tag2;
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
  multicast {
    interface interface-name {
      no-qos-adjust;
    }
  }
  rib routing-table-name {
    access {
      route prefix {
        metric route-cost;
        next-hop next-hop;
        preference route-distance;
        tag route-tag;
        tag2 route-tag2;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
  }
}

```

Hierarchy Level

```
[edit dynamic-profiles profile-name],  
[edit dynamic-profiles profile-name routing-instances $junos-routing-instance]
```

Release Information

Statement introduced in Junos OS Release 9.6.

Support at the `[edit dynamic-profiles profile-name routing-instances $junos-routing-instance]` hierarchy level introduced in Junos OS Release 10.1.

Description

Configure protocol-independent routing properties in a dynamic client profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Dynamic Access Routes for Subscriber Management

Configuring Dynamic Access-Internal Routes for DHCP and PPP Subscribers

routing-instance (Subscriber Secure Policy)

Syntax

```
routing-instance routing-instance-name;
```

Hierarchy Level

```
[edit services radius-flow-tap]
```

Release Information

Statement introduced in Junos OS Release 15.1R3 for enhanced subscriber management on MX Series routers.

Description

Specify the routing instance that is used to send mirrored packets to a mediation device for subscriber secure policy traffic mirroring.

Options

routing-instance-name—Name of the routing instance.

Default: Routing instance **default**

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Subscriber Secure Policy Overview | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

[Configuring Support for Subscriber Secure Policy Mirroring | 574](#)

routing-instance (PCC Action Profiles)

Syntax

```
routing-instance {
  downlink downlink-vrf-name;
  uplink uplink-vrf-name;
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name steering],
[edit services pcef pcc-action-profiles profile-name steering]
```

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name* steering]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the routing instance that a PCC action profile uses for steering traffic.

Options

downlink *downlink-vrf-name*—Use the specified name of the routing instance for downlink traffic (to the access side) or the predefined dynamic interface variable .

uplink *uplink-vrf-name*—Use the specified name of the routing instance for uplink traffic (from the access side).

NOTE: The routing instances must have been previously configured.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

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

Reduce forwarding of IP packets that might be spoofing and address by checking whether traffic is arriving on an expected path that the sender would use to reach the destination. You can include this statement with the **inet** protocol family only. When the traffic passes the check, it is forwarded to the destination address; otherwise it is discarded. When you configure **rpf-check** alone, then unicast RPF is in strict mode, meaning that the check passes only when the packet's source address is in the FIB and the interface matches the routes RPF.

Starting in Junos OS Release 19.1, the **show interfaces statistics logical-interface-name detail** command displays unicast RPF statistics for dynamic logical interfaces when either **rpf-check** or **rpf-check mode loose** is enabled on the interface. No additional statistics are displayed when **rpf-check fail-filter filter-name** is configured on the interface. The **clear interfaces statistics logical-interface-name** command clears RPF statistics.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Unicast RPF in Dynamic Profiles for Subscriber Interfaces](#) | 347

[Configuring Unicast RPF Strict Mode](#)

rule (Captive Portal Content Delivery)

Syntax

```
rule rule-name {
  match-direction (input | output | input-output);
  term term-name {
    from {
      destination-address address <except>;
    }
    then {
      accept;
      insert tag-name tag-name tag-value tag-value;
      redirect url;
      rewrite destination-address address <destination-port port-number>;
      syslog;
    }
  }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery],
[edit services captive-portal-content-delivery]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the **[edit dynamic-profiles profile-name services captive-portal-content-delivery rule rule-name term term-name]** hierarchy level added in Junos OS Release 17.2R1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

Description

Specify the rule the router uses when applying this service. Use the statement at the **[edit services...]** hierarchy level for static CPCD. Use the statement at the **[edit dynamic-profiles profile-name services...]** hierarchy level for converged services CPCD.

Options

rule-name—Identifier for the collection of terms that constitute this rule.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview | 473](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

rule (LRF)

Syntax

```
rule lrf-rule-name {
  then {
    report {
      collector collector-name;
      template template-name;
      time-limit time-interval;
      volume-limit volume;
    }
  }
}
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure an LRF rule, which controls how data sessions are logged and reported. In this release, the matching conditions for an LRF rule are identified in a static PCC rule, not in the LRF rule.

Options

lrf-rule-name—Name of the LRF rule.

Range: Up to 63 characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers](#) | 460

[Configuring Logging and Reporting for Junos OS Subscriber Aware](#)

rule-set (Captive Portal Content Delivery)

Syntax

```
rule-set rule-set-name {  
    [rule rule-name];  
}
```

Hierarchy Level

```
[edit services captive-portal-content-delivery]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

Description

Define a set of captive portal content delivery rules that the router uses when applying this service. This statement is supported only for static CPCD; it is not supported for converged services CPCD.

Options

rule-set-name—Identifier for the collection of rules that constitute this rule set.

rule rule-name—Name of a rule defined at the `[edit services captive-portal-content-delivery]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview | 473](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

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](#) | 40

[Configuring Schedulers in a Dynamic Profile for Subscriber Access](#) | 52

[Dynamic Variables Overview](#)

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 | 40](#)

[Configuring Traffic Scheduling and Shaping for Subscriber Access | 47](#)

[output-traffic-control-profile | 1006](#)

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 client profile or a dynamic service 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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access](#) | 40

[Configuring Schedulers in a Dynamic Profile for Subscriber Access](#) | 52

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 client profile or a dynamic service 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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)[Configuring Schedulers in a Dynamic Profile for Subscriber Access | 52](#)[scheduler | 1134](#)

service (Dynamic Profiles)

Syntax

```
service {  
  pcef pcef-profile-name {  
    activate rule-name | activate-all;  
  }  
}
```

Hierarchy Level

```
[edit dynamic-profile profile-name interfaces interface-name unit logical-unit-number]
```

Release Information

Statement introduced in Junos OS Release 17.2 on MX Series.

Description

Assign a policy and charging enforcement function (PCEF profile) to the dynamic profile. The PCEF profile specifies a set of PCC rules and rulebases to assign to a subscriber for application-aware policy control, and assigns a precedence value to each predefined rule. Also specify which of the PCC rules to activate for the dynamic profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Application-Aware Policy Control for Subscriber Management | 410](#)

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Dynamic Service Sets Overview](#) | 365

Associating Service Sets with Interfaces in a Dynamic Profile | 366

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494

Configuring Routing Engine-Based, Static HTTP Redirect Services | 505

Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519

service-agents (Analytics)

Syntax

```

service-agents {
  agent-name {
    inputs {
      analytics {
        parameters {
          generate-tags value;
          sample-frequency value;
          sensors file-path;
        }
      }
      input-ipfix {
        parameters {
          maximum-connections number;
          tcp-port port-number;
          vrf-name name;
        }
      }
      input-jti-ipfix {
        parameters {
          record-group group-name {
            record ipfix-record-name;
            reporting-interval seconds;
          }
        }
      }
    }
    outputs {
      file {
        parameters {
          path file-path;
        }
      }
      kafka {
        parameters {
          server ip-address;
          topic topic-name;
          encoding encoding-type;
        }
      }
      output-ipfix {
        parameters {

```

```

        collector-address ip-address;
        collector-ca-certificate file-path;
        collector-certificate file-path;
        collector-certificate-key file-path;
        collector-connection-retry-interval seconds;
        collector-tcp-port port-number;
        collector-vrf-name vrf-name;
    }
}
}
}
}

```

Hierarchy Level

[edit services analytics [agent](#)]

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Statement introduced in Junos OS Release 18.4R1 on PTX Series routers.

Description

Configure a network analytics service agent that uses input and output plug-ins to collect, transform, and forward network telemetry data.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system

RELATED DOCUMENTATION

[Configuring the BNG as an IPFIX Mediator to Collect and Export IPFIX Data](#) | 672

[Configuring the Collection and Export of Local Telemetry Data on the IPFIX Mediator](#) | 680

[Configuring NTF Agent](#)

[IPFIX Mediation on the BNG](#) | 667

[Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector](#) | 677

service-device (Remote Device Management)

Syntax

```

service-device device-name {
  access-domain {
    vlan-id-list [vlan-id-low-vlan-id-high vlan-id]
  }
  address ip-address;
  dictionary absolute file path;
  provisioning-method {
    netconf {
      bulk-interval milliseconds;
      bulk-limit number;
      connection-retry-interval seconds;
      password password;
      port port-number;
      reconfigure-bulk-limit number;
      response-timeout seconds;
      response-timeout-count number;
      user-name name;
    }
  }
}

```

Hierarchy Level

```

[edit routing-instances routing-instance-name system services remote-device-management],
[edit system services remote-device-management]

```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Configure a remote device to provision and deprovision services for subscribers. This configuration is used when **profile-type remote-device-service** has been configured in the dynamic service profile.

Provides basic configuration for the remote service device and enables NETCONF TCP session to be established with the service-device. Subsequently, dynamic configuration of the access-domain occurs based on allocation and assignment of individual VLAN IDs or VLAN ranges for one or more subscriber sessions.

NOTE: With the exception of the **access-domain** statement, all statements are required to be configured for the service device and are subject to a commit check. This behavior enables basic configuration for the remote device and the NETCONF TCP session to be completed followed later by dynamic configuration of the access-domain based on allocation and assignment of individual VLAN IDs or VLAN ranges for one or more subscriber sessions.

Options

address *ip-address*—Specify the IP address of the remote device used by the BNG to configure the subscriber service. The address must be unique; it is used for all actions performed by the BNG, including service provisioning and deprovisioning. The address must also be unique across all routing instances.

NOTE: You cannot change the IP address for the device when any active subscriber services are mapped to it.

device-name—System-wide name that uniquely identifies the device across routing instances.

dictionary *absolute file path*—Specify the absolute file path for the vendor-specific dictionary that defines the set of NETCONF XML protocol commands required to provision, deprovision, and roll back a subscriber service on the remote device. The dictionary is stored on the BNG. An example absolute path is **/var/home/dict/remote-device.xml**.

Because the dictionary is assigned to a service device, you can use service devices from different vendors at the same time. This assignment method also enables you to use multiple devices from the same vendor that are running different software releases. This is useful because the NETCONF XML protocol commands to configure the devices can differ across releases.

NOTE: The maximum total length of the path is 127 characters. The filename must end in **xml**.

NOTE: You cannot change the dictionary path for the device when any active subscriber services are mapped to it.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

RELATED DOCUMENTATION

[Configuring Remote Device Management for Service Provisioning | 648](#)

[Remote Device Services Manager \(RDSM\) Overview | 631](#)

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. You can use the predefined dynamic interface variables **\$junos-input-service-filter**, **\$junos-output-service-filter**, **\$junos-input-ipv6-service-filter**, and **\$junos-output-ipv6-service-filter**. 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.

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](#) | 365

[Associating Service Sets with Interfaces in a Dynamic Profile | 366](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

service-interface (Services Interfaces)

Syntax

```
service-interface interface-name;
```

Hierarchy Level

```
[edit services service-set service-set-name interface-service]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the name for the services interface associated with an interface-wide service set.

Options

interface-name—Identifier 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 Service Sets to be Applied to Services Interfaces](#)

[Applying Services to Subscriber-Aware Traffic with a Service Set](#)

service-set (Application-Aware Control Policy)

Syntax

```
service-set service-set-name {
  application-identification-profile app-id-profile-name;
  interface-service {
    service-interface interface-name;
  }
  lrf-profile profile-name;
  pcef-profile pcef-profile-name;
  service-set-options {
    subscriber-awareness;
  }
}
```

Hierarchy Level

```
[edit services]
```

Release Information

Statement options for application-aware policy control introduced in Junos OS Release 17.2R1 on MX Series for Broadband Subscriber Management.

USF support for statement options for application-aware policy control introduced in Junos OS Release 19.3R2 on MX Series for Broadband Subscriber Management.

Description

Configure a service set to identify the service interface that handles application-aware policy control.

Options

application-identification-profile *app-id-profile-name*—Specify a dummy application identification profile that you configured at the **[edit services application-identification profile]** hierarchy level. This profile is a placeholder profile with no configuration options, but it must be specified to enable application identification functionality on the services plane.

lrf-profile *profile-name*—Identify the LRF profile that defines the logging of subscriber application-aware data sessions. The LRF profile must first be configured at the **[edit services lrf]** hierarchy level.

pcef-profile *pcef-profile-name*—Specify a dummy PCEF profile that you configured at the **[edit services pcef]** hierarchy level. This profile is a placeholder profile with no configuration options, but it must be specified to enable application-aware policy control functionality on the services plane.

service-interface *interface-name*—Specify the services PIC interface on which the services are performed.

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

Identifying the Service Interface That Handles Subscriber Management Application-Aware Policy Control | 424

Configuring Logging and Reporting for Subscriber Management | 459

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.

From 17.2R1 onwards, you can configure converged services at the **edit dynamic-profiles http-redirect-converged** hierarchy level.

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. You can use the predefined dynamic interface variables **\$junos-input-service-set**, **\$junos-output-service-set**, **\$junos-input-ipv6-service-set**, and **\$junos-output-ipv6-service-set**.

NOTE: Starting in Junos OS Release 17.2R1, you can configure converged services at the **edit dynamic-profiles http-redirect-converged** hierarchy level. CPCD rules can also be configured under the dynamic profiles stanza to achieve parameterization of the rules. This mechanism provides additional flexibility to customize the different rules on a per subscriber basis through service attachment.

Options

service-set-name—Name of the service set.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Dynamic Service Sets Overview | 365](#)

[Associating Service Sets with Interfaces in a Dynamic Profile | 366](#)

services (Captive Portal Content Delivery)

Syntax

```

services {
  ...
  captive-portal-content-delivery {
    auto-deactivate value;
    profile name
      cpcd-rule-sets rule-set-name;
      cpcd-rules rule-name;
      dynamic;
      http-redirect-options url;
      ipda-rewrite-options {
        destination-address destination-address;
        destination-port destination-port;
      }
    }
  rule rule-name {
    match-direction (input | output | input-output);
    from {
      destination-address address <except>;
    }
    term term-name {
      then {
        accept;
        insert tag tag-name tag-value tag-value;
        redirect url;
        rewrite destination-address address <destination-port port-number>;
        syslog;
      }
    }
  }
  rule-set rule-set-name {
    [rule rule-name];
  }
  traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-readable)>;
    flag name;
    no-remote-trace no-remote-trace;
  }
}

```

Hierarchy Level

```
[edit],
[edit dynamic-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the **[edit dynamic-profiles *profile-name* services]** hierarchy level added in Junos OS Release 17.2R1.

Description

Define the captive portal content delivery set of the rules statements to be applied to traffic. Use the statement at the **[edit services...]** hierarchy level for static CPCD. Use the statement at the **[edit dynamic-profiles *profile-name* services...]** hierarchy level for converged services CPCD.

The **profile**, **rule-set**, and **traceoptions** stanzas are not supported at the **[edit dynamic-profiles *profile-name* hierarchy level]**.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview | 473](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

session-options

Syntax

```
session-options {
  client-group [ group-names ];
  client-idle-timeout minutes;
  client-idle-timeout-ingress-only;
  client-session-timeout minutes;
  pcc-context {
    input-service-filter-name filter-name;
    input-service-set-name service-set-name;
    ipv6-input-service-filter-name filter-name;
    ipv6-input-service-set-name service-set-name;
    ipv6-output-service-filter-name filter-name;
    ipv6-output-service-set-name service-set-name;
    output-service-filter-name filter-name;
    output-service-set-name service-set-name;
    profile-name pcef-profile-name;
  }
  strip-user-name {
    delimiter [ delimiter ];
    parse-direction (left-to-right | right-to-left);
  }
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 8.5.

Description

(MX Series and SRX Series devices) Define options to place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both.

(MX Series) Define options to modify a subscriber username at login based on the subscriber's access profile.

(MX Series) Specify characteristics related to policy and charging control (PCC) rules, such as the PCEF profile that contains the rules, service sets to process the rules, and service filters for the service sets.

Options

client-idle-timeout—(MX Series only) Specify the grace period that begins after an authenticated user terminates all sessions and connections. Authentication is not required if a new connection is initiated during the grace period by the same user.

During this period, the router determines whether the subscriber is inactive by monitoring data traffic, both upstream from the user (ingress) and downstream to the user (egress). Control traffic is ignored. The subscriber is not considered idle as long as data traffic is detected in either direction. When no traffic is detected for the duration of the idle time out, non-DHCP subscribers (such as L2TP or PPP) are gracefully logged out, similarly to a RADIUS-initiated disconnect or a CLI-initiated logout; DHCP subscribers are disconnected.

When you additionally configure the related **client-idle-timeout-ingress-only** statement (MX Series only), the router monitors only ingress traffic to determine whether the subscriber is inactive; it does not monitor any egress traffic. The related **client-session-timeout** statement terminates the subscriber session when the session timeout expires regardless of user activity.

Client idle timeouts are most often used for residential services rather than business services. The most practical use case for this timeout is in a PPP access model. It is not practical for DHCP or DHCPv6 subscribers.

Although you can use the **client-idle-timeout** statement for dynamically configured subscriber VLANs, this configuration is useful only in limited circumstances (such as IP over Ethernet without DHCP and with fixed addresses) and is not typically used. If you do use the idle timeout for VLANs, the timeout period starts when the VLAN is instantiated. It resets when a client session is created or an existing session is reactivated. When no traffic is detected on an authenticated VLAN for the duration of the timeout, the VLAN is considered inactive and is deleted. If no client sessions are ever created on the VLAN, then the VLAN is removed when the timeout expires.

Default: The timeout is not configured.

Values: *minutes*—Number of minutes of idle time that elapse before the session is terminated. The value that you specify must be determined locally with consideration of the services and policies that you offer.

Range: 10 through 1440 minutes

client-idle-timeout-ingress-only—(MX Series only) Starting in Junos OS Release 16.2, specify that only ingress traffic is monitored for subscriber idle timeout processing for the duration of the idle timeout period that you specify with the **client-idle-timeout** statement. If no ingress traffic is received for the duration of the timeout, then the subscriber is gracefully logged out (non-DHCP subscribers) or disconnected (DHCP subscribers).

If you configure **client-idle-timeout** alone, then both ingress and egress traffic are monitored during the idle timeout. Monitoring only ingress traffic is useful in cases where the LNS sends traffic to the remote peer even when the peer is not up, such as when the LNS does not have PPP keepalives enabled and therefore does not detect that the peer is not up. Because the LAC monitors both ingress and egress traffic by default, in this situation it receives the egress traffic from the LNS and either

does not log out the subscriber or delays detection of inactivity until the egress traffic ceases. When you specify that only ingress traffic is monitored in this case, the LAC can detect that the peer is inactive and then initiate logout.

client-session-timeout—(SRX Series, vSRX only) Specify the amount of time after which user sessions are terminated, regardless of user activity (also known as a forced or hard authentication timeout).

Alternatively, when you want subscribers to be identified as inactive before they are terminated, use the related statements, **client-idle-timeout** and **client-idle-timeout-ingress-only**. Use **client-idle-timeout** alone to specify a period of time during which both ingress and egress subscriber data traffic is monitored; if no traffic is detected for the duration of the period, the subscriber is considered inactive and is terminated. Add the **client-idle-timeout-ingress-only** statement to monitor only ingress traffic for the duration of the timeout set with the **client-idle-timeout** statement.

BEST PRACTICE: We recommend that you do not configure a session timeout for subscribers receiving voice services. Because the session timeout is a simple time-based timeout, it is likely to interrupt subscribers actively using a voice service and terminate their calls unexpectedly (from the subscriber viewpoint). This result is a particular concern for emergency services calls.

Client session timeouts are most often used for residential services rather than business services. The most practical use case for this timeout is in a PPP access model when no voice services are offered. For DHCP or DHCPv6 subscribers, the session timeout is used as the DHCP lease timer if no other lease time configuration is present.

Although you can use the **client-session-timeout** statement for dynamically configured subscriber VLANs, this configuration is useful only in limited circumstances (such as IP over Ethernet without DHCP and with fixed addresses) and is not typically used. If you do use the session timeout for VLANs, the timeout period starts when the VLAN is instantiated.

Default: The timeout is not configured.

Values: *minutes*—Number of minutes after which user sessions are terminated. The value that you specify must be determined locally with consideration of the services and policies that you offer.

Range: 1 through 527040 minutes

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Subscriber Session Timeout Options

Configuring Username Modification for Subscriber Sessions

Removing Inactive Dynamic Subscriber VLANs

[Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management](#) | 428

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 | 40](#)

[Configuring Traffic Scheduling and Shaping for Subscriber Access | 47](#)

[output-traffic-control-profile | 1006](#)

shared-name

Syntax

```
shared-name filter-shared-name;
```

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 interface-name unit logical-unit-number filter input filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number filter 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

Guidelines for Configuring Firewall Filters

[Understanding Dynamic Firewall Filters | 226](#)

[Classic Filters Overview | 229](#)

[Basic Classic Filter Syntax | 232](#)

signature (Application Identification)

Syntax

```
signature l4-l7-signature-name {
  chain-order
  member member-name {
    check-bytes max-bytes-to-check;
    context context;
    pattern pattern;
    direction (any | client-to-server | server-to-client);
  }
  order order;
  order-priority (high | low);
  port-range {
    tcp [port-range];
    udp [port-range];
  }
  protocol (http | ssl | tcp | udp);
```

Hierarchy Level

[edit services application-identification application *application-name* over *protocol-type*]

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure an application signature for pattern matching.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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](#) | [784](#)

[color-blind](#) | [786](#)

[two-rate](#) | [1229](#)

snmp (Subscriber Secure Policy)

Syntax

```
snmp {  
    notify-targets ip-address;  
}
```

Hierarchy Level

[edit services [radius-flow-tap](#)]

Release Information

Statement introduced in Junos OS Release 16.1R1.

Description

Specify the IP address for a target mediation device to receive SNMPv3 encrypted trap notifications. Only these configured targets can receive the notifications. This is required for secure SNMPv3 notifications for subscriber secure policy mirroring.

Options

notify-targets *ip-address*—Specify the IP address of a trap target that is allowed to receive encrypted SNMPv3 subscriber secure policy mirroring traps. If you configure multiple targets, you must configure them one at a time.

Required Privilege Level

flow-tap—To view this statement in the configuration.

flow-tap-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Support for Subscriber Secure Policy Mirroring](#) | [574](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview](#) | [572](#)

[Subscriber Secure Policy Overview](#) | [553](#)

source (Application Identification)

Syntax

```
source ip ip-address-prefix;
```

Hierarchy Level

```
[edit services application-identification application application-name address-mapping]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Specify the source IP address for address mapping-based application identification.

Options

ip-address-prefix—IP address and prefix for matching.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

[Application Identification Overview | 431](#)

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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

Enabling IGMP Static Group Membership

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

[Dynamic MLD Configuration Overview](#) | 401

Enabling MLD Static Group Membership

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

source-address (LRF Profile)

Syntax

```
source-address source-address;
```

Hierarchy Level

```
[edit services lrf profile profile-name collector collector-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure the source address to be used when exporting data to the collector.

Options

source-address—IP address to be used as the source address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers](#) | 460

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management](#) | 459

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

[Dynamic MLD Configuration Overview](#) | 401

Enabling MLD Static Group Membership

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

[Dynamic MLD Configuration Overview](#) | **401**

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

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 | 553](#)

[Configuring RADIUS-Initiated Subscriber Secure Policy Mirroring Overview | 572](#)

ssh (System Services)

Syntax

```
ssh {
  authentication-order [method 1 method2...];
  ciphers [ cipher-1 cipher-2 cipher-3 ...];
  client-alive-count-max seconds;
  client-alive-interval seconds;
  connection-limit limit;
  fingerprint-hash (md5 | sha2-256);
  hostkey-algorithm (algorithm | no-algorithm);
  key-exchange [algorithm1 algorithm2...];
  log-key-changes log-key-changes;
  macs [algorithm1 algorithm2...];
  max-pre-authentication-packets number;
  max-sessions-per-connection number;
  no-passwords;
  no-public-keys;
  no-tcp-forwarding;
  protocol-version [v2];
  rate-limit limit;
  root-login (allow | deny | deny-password);
  sftp-server;
}
tcp-forwarding
```

Hierarchy Level

```
[edit system services]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

client-alive-interval and **client-alive-max-count** statements introduced in Junos OS Release 12.2.

no-passwords statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

no-public-keys statement introduced in Junos OS release 15.1.

tcp-forwarding statement introduced in Junos OS Release 15.1X53-D50 for the NFX250 Network Services Platform.

fingerprint-hash statement introduced in Junos OS Release 16.1.

log-key-changes statement introduced in Junos OS Release 17.4R1.

sftp-server statement introduced in Junos OS Release 19.1R1.

Description

Allow SSH requests from remote systems to access the local router or switch.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring SSH Service for Remote Access to the Router or Switch*

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.

Statement became non-functional in Junos OS Release 15.1R4.

Statement deprecated in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1.

Statement removed from CLI in Junos OS Release 17.3R1.

Description

Apply an SSM map to a dynamic IGMP interface. SSM mapping translates IGMPv1 and IGMPv2 membership reports to an IGMPv3 report, which enables hosts running IGMPv1 or IGMPv2 to participate in SSM. The SSM map associates an SSM policy that matches group addresses to be translated with the source addresses where the group addresses are found.

Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the **ssm-map** statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the statement is removed from the CLI. Use the **ssm-map-policy** instead to associate the policy with the dynamic IGMP interface for all releases with enhanced subscriber management.

If you upgrade from an earlier release that does not support enhanced subscriber management (any release earlier than Junos OS Release 15.1R4) with a configuration that includes the **ssm-map** statement, the results vary depending on the release to which you are upgrading:

- Upgrade to Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, or later releases in those branches—The **ssm-map** configuration is allowed and does not cause the upgrade to fail. However, the configuration has no effect and subscribers cannot log in.
- Upgrade to Junos OS Release 17.3R1 or later releases—The upgrade fails because the **ssm-map** configuration is not allowed. If you perform the upgrade without validation (**no-validate**), the upgrade passes and the **ssm-map** configuration is accepted, but it has no effect.

BEST PRACTICE: Delete the **ssm-map** configuration before you upgrade.

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.

Release History Table

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the ssm-map statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the statement is removed from the CLI.

RELATED DOCUMENTATION

Dynamic IGMP Configuration Overview 393
Configuring Dynamic DHCP Client Access to a Multicast Network 395
<i>Source-Specific Multicast Groups Overview</i>

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.

Statement became non-functional in Junos OS Release 15.1R4.

Statement deprecated in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1.

Statement removed from CLI in Junos OS Release 17.3R1.

Description

Apply an SSM map to a dynamic MLD interface. SSM mapping translates MLDv1 membership reports to an MLDv2 report, which enables hosts running MLDv1 to participate in SSM. The SSM map associates an SSM policy that matches group addresses to be translated with the source addresses where the group addresses are found.

Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the **ssm-map** statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the statement is removed from the CLI. Use the **ssm-map-policy** instead to associate the policy with the dynamic MLD interface for all releases with enhanced subscriber management.

If you upgrade from an earlier release that does not support enhanced subscriber management (any release earlier than Junos OS Release 15.1R4) with a configuration that includes the **ssm-map** statement, the results vary depending on the release to which you are upgrading:

- Upgrade to Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, or later releases in those branches—The **ssm-map** configuration is allowed and does not cause the upgrade to fail. However, the configuration has no effect and subscribers cannot log in.
- Upgrade to Junos OS Release 17.3R1 or later releases—The upgrade fails because the **ssm-map** configuration is not allowed. If you perform the upgrade without validation (**no-validate**), the upgrade passes and the **ssm-map** configuration is accepted, but it has no effect.

BEST PRACTICE: Delete the **ssm-map** configuration before you upgrade.

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.

Release History Table

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the ssm-map statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the statement is removed from the CLI.

RELATED DOCUMENTATION

Dynamic MLD Configuration Overview 401
<i>Example: Configuring SSM Mapping</i>

ssm-map-policy (Dynamic IGMP Interface)

Syntax

```
ssm-map-policy ssm-map-policy-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols igmp interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 16.1R5.

Description

Apply an SSM map policy to a dynamic IGMP interface. SSM mapping translates IGMPv1 and IGMPv2 membership reports to an IGMPv3 report, which enables hosts running IGMPv1 or IGMPv2 to participate in SSM. The map policy associates the group addresses to be translated with the source addresses where the group addresses are found. You configure the SSM map policy with the **policy-statement** statement at the **[edit policy-options]** hierarchy level.

For statically-configured IGMP interfaces, use the **ssm-map-policy (IGMP)** statement.

Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1, use this statement instead of the **ssm-map** statement to associate the policy with the dynamic IGMP interface for all releases with enhanced subscriber management. Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the **ssm-map** statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the statement is removed from the CLI.

Required Privilege Level

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

Release History Table

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1, use this statement instead of the ssm-map statement to associate the policy with the dynamic IGMP interface for all releases with enhanced subscriber management.

RELATED DOCUMENTATION

[Configuring SSM Mapping for Dynamic IGMP and MLD | 399](#)

[Dynamic IGMP Configuration Overview | 393](#)

Example: Configuring SSM Maps for Different Groups to Different Sources

ssm-map-policy (Dynamic MLD Interface)

Syntax

```
ssm-map-policy ssm-map-policy-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name protocols mld interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 16.1R5.

Description

Apply an SSM map policy to a dynamic MLD interface. SSM mapping translates MLDv1 membership reports to an MLDv2 report, which enables hosts running MLDv1 to participate in SSM. The map policy associates the group addresses to be translated with the source addresses where the group addresses are found. You configure the SSM map policy with the **policy-statement** statement at the **[edit policy-options]** hierarchy level.

For statically-configured MLD interfaces, use the **ssm-map-policy (MLD)** statement.

Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1, use this statement instead of the **ssm-map** statement to associate the policy with the dynamic MLD interface for all releases with enhanced subscriber management. Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, and 17.2R2, the **ssm-map** statement is deprecated and no longer supported. Starting in Junos OS Release 17.3R1, the **ssm-map** statement is removed from the CLI.

Required Privilege Level

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

Release History Table

Release	Description
16.1R5	Starting in Junos OS Release 16.1R5, 16.2R3, 17.1R3, 17.2R2, and 17.3R1, use this statement instead of the ssm-map statement to associate the policy with the dynamic MLD interface for all releases with enhanced subscriber management.

RELATED DOCUMENTATION

[Configuring SSM Mapping for Dynamic IGMP and MLD | 399](#)

[Dynamic MLD Configuration Overview | 401](#)

Example: Configuring SSM Maps for Different Groups to Different Sources

stacked-interface-set (Dynamic Profiles)

Syntax

```
stacked-interface-set {
  interface-set-name interface-set-name {
    interface-set-name interface-set-name;
  }
}
```

Hierarchy Level

[edit [dynamic-profiles](#) *profile-name* [interfaces](#)]

Release Information

Statement introduced in Junos OS Release 18.4R1.

Description

For MX Series routers with MPCs that support 5-level hierarchical scheduling, define an interface set over interface set hierarchy. The child interface set (L3) contains subscriber logical interfaces, and the parent interface set (L2) contains one or more interface sets as members.

NOTE: For both interface-sets in the stack (at L3 and L2), a CoS traffic-control-profile (TCP) must be assigned to each.

Options

interface-set-name—Name of the interface set to be configured or one of the following Junos OS predefined variables:

- **\$junos-aggregation-interface-set-name**—L2 interface-set representing a logical intermediate node (DPU-C or PON tree) in the access network.
- **\$junos-default-interface-set-name**—An intermediate, default variable used to resolve the L3 (child) interface-set name based on whether subscriber access is via a logical intermediate aggregation node
- **\$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-phy-ifd-underlying-intf-set-name**—A default, topology-based interface-set derived from the predefined variable, \$junos-phy-ifd-underlying-intf-set-name by appending “-underlying” to conserve

L2 CoS nodes. It is a variation of the `$junos-phy-ifd-interface-set-name` variable used as a default L2 interface-set for 4-level hierarchy or default L3 interface-set for 5-level hierarchy.

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](#) | 189

[interface-set \(Dynamic Profiles\)](#) | 946

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

[Enabling IGMP Static Group Membership](#)

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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

[Dynamic MLD Configuration Overview](#) | 401

Enabling MLD Static Group Membership

static-policy-control

Syntax

```
static-policy-control {
  pcc-rules {
    [rule-name precedence number <time-of-day-profile profile-name>];
  }
  pcc-rulebases {
    [rulebase-name <time-of-day-profile profile-name>];
  }
}
```

Hierarchy Level

```
[edit unified-edge pcef profiles profile-name],
[edit services pcef profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef profiles profile-name]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Description

Configure static policy control for the policy and charging control (PCC) rules or PCC rulebase in a policy and charging enforcement function (PCEF) profile. You can configure a maximum of 32 PCC rules in a PCEF profile. There is no limit to the number of PCC rulebases you can configure in a PCEF profile.

NOTE: For Junos OS Subscriber Aware, you can configure only one of the following statements in a PCEF profile: **aaa-policy-control**, **static-policy-control**, or **dynamic-policy-control**. For Junos OS Subscriber Management, you can configure only **static-policy-control**.

If you are using Junos OS Subscriber Aware, configure static policy control at the **[edit unified-edge pcef profiles profile-name]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure static policy control at the **[edit services pcef profiles profile-name]** hierarchy level.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring a Policy and Charging Enforcement Function Profile for Junos OS Subscriber Aware Static Policies

[Configuring a Policy and Charging Enforcement Function Profile for Subscriber Management](#) | 423

steering

Syntax

```
steering {
  keep-existing-steering;
  path {
    ipv4-address ipv4-address;
    ipv6-address ipv6-address;
  }
  routing-instance {
    downlink downlink-vrf-name;
    uplink uplink-vrf-name;
  }
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name],
[edit services pcef pcc-action-profiles profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the method that a PCC action profile uses for steering traffic

If you are using Junos OS Subscriber Aware, configure steering at the **[edit unified-edge pcef pcc-action-profiles *profile-name*]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, configure the PCC action profile at the **[edit services pcef pcc-action-profiles *profile-name*]** hierarchy level.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Application-Aware Policy Control for Subscriber Management | 406](#)

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

Understanding Predefined Policy and Charging Control Rules for Subscriber-Aware Traffic Treatment

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

subscriber-leave-timer

Syntax

```
subscriber-leave-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast
  interface interface-name],
[edit logical-systems logical-system-name routing-options multicast interface interface-name],
[edit routing-instances routing-instance-name routing-options multicast interface interface-name],
[edit routing-options multicast interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 9.2.

Statement introduced in Junos OS Release 9.2 for EX Series switches.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.

Options

seconds—Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update. This is the same as if the statement were not configured.

Range: 0 through 30

Default: 0 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

tags (Application Identification)

Syntax

```
tags tag-name tag-value;
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Specify an application tag that provides general information about the application, such as associated risk factors, technology, and the type of traffic. The tag consists of a user-defined name and value.

Options

tag-name—Name for the tag, which is a textual string.

tag-value—Value for the tag.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

[Application Identification Overview | 431](#)[Application Identification Overview | 431](#)

targeted-distribution (Dynamic Demux Interfaces over Aggregated Ethernet)

Syntax

```
targeted-distribution;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Description

Configure egress data for a dynamic logical interface to be sent across a single member link in an aggregated Ethernet bundle. A backup link is provisioned and CoS scheduling resources are switched to the backup link 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

| *Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces*

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.

Statement introduced in Junos OS Release 13.2R2 for EX Series switches.

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 and CoS scheduling resources are switched to the backup link 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 | 45](#)

Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces

Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface

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

tcp-forwarding (Processes)

Syntax

```
tcp-forwarding {
  disable;
  traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
```

Hierarchy Level

[edit system processes]

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Disable the TCP port forwarding process or configure tracing operations for TCP port forwarding events.

Options

disable—Disable the TCP port forwarding process, tcpfwdd.

The remaining statement is explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Tracing TCP Port Forwarding Events for Troubleshooting | 662](#)

[TCP Port Forwarding for Remote Device Management | 655](#)

tcp-forwarding (Remote Device Management)

Syntax

```
tcp-forwarding {
  listening-port port-number listening-address ipv4-listening-address {
    allowed-source ipv4-prefix;
    forwarding-address ipv4-forwarding-address;
    forwarding-port forwarding-port-number;
    max-connections number;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system services],
[edit logical-systems logical-system-name system services],
[edit routing-instances routing-instance-name system services],
[edit system services]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Configure the mapping between the TCP listening address/listening port combination on the BNG and the TCP port forwarding address/port combination where the BNG forwards the incoming data stream. TCP port forwarding is used when the BNG, together with one or more access nodes, is treated as a single addressable point of management by an external management system. The TCP port forwarding connections enable the BNG to demultiplex and multiplex management requests exchanged between the access nodes and the management system.

Options

allowed-source *ipv4-prefix*—(Optional) Restrict the IPv4 prefixes from which TCP connections are accepted on the listening port. The **allowed-source** value is compared to the source address in the TCP header from the triggering entity. When you do not configure an allowed source, TCP connections are accepted from any source prefix.

You can use a /32 IPv4 mask to specify a single address as the source or you can use other masks to specify an IPv4 subnet as the source. You can configure an unlimited number of prefixes for each listening port. To configure multiple sources, you must include the statement multiple times, once for each additional source prefix.

NOTE: You can also configure an unlimited number of allowed-source prefixes across the system.

forwarding-address *ipv4-forwarding-address*—Specify the IPv4 address to which MX BNG must open the second connection of the TCP pair after it opens the first connection triggered on the listening port/listening address combination. All packets received on one connection of the TCP pair are transmitted on the peer (second) connection. This address is used with the forwarding port to open the peer connection.

forwarding-port *forwarding-port-number*—Specify the TCP port of the peer (second) connection of the TCP pair. This port is used with the forwarding address to open the peer connection.

Range: 1 through 65,535

listening-address *ipv4-listening-address*—Specify a particular IPv4 address on the BNG that a triggering entity (an external management or provisioning system or a remote device) must use when attempting to trigger connections on the listening port. You must configure a unique combination of listening port and listening address for each TCP mapping.

listening-port *port-number*—Specify the TCP port that the BNG monitors for connections to be triggered by a remote device or an external management or provisioning system.

Range: 8000 through 8031

max-connections *number*—(Optional) Set a limit on the number of simultaneous TCP connections that the BNG allows on a single listening port. Connection requests received after this limit is reached are rejected.

NOTE: In addition to this per-listening port limit, the system-wide limit for TCP connections is 128 (64 pairs) across all routing instances and listening ports.

Range: 1 through 16

Default: 1

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [TCP Port Forwarding for Remote Device Management](#) | 655

template (LRF Profile)

Syntax

```
template template-name {
  format ipfix;
  template-tx-interval tx-time;
  template-type template-type;
  trigger-type (session-close | volume);
}
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure a template, which specifies a set of data to be transmitted. This template can be specified in LRF rules.

Options

template-name—Name for the template.

Range: Up to 32 characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers](#) | 460

[Configuring Logging and Reporting for Junos OS Subscriber Aware](#)

[Configuring Logging and Reporting for Subscriber Management](#) | 459

template (LRF Rule)

Syntax

```
template template-name;
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name then report]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Specify the template that identifies the type of data to report if the LRF rule is matched.

Options

template-name—Name of the template that is used. The referenced template must be configured.

Range: Up to 32 characters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

[Configuring Logging and Reporting for Junos OS Subscriber Aware](#)

[Configuring Logging and Reporting for Subscriber Management | 459](#)

template-tx-interval (LRF Profile)

Syntax

```
template-tx-interval tx-time;
```

Hierarchy Level

```
[edit services lrf profile profile-name template template-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure the interval at which to retransmit the template to the collector.

Options

tx-time—Time interval in seconds.

Default: 60

Range: 10 through 600

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management | 459](#)

template-type (LRF Profile)

Syntax

```
template-type template-type;
```

Hierarchy Level

```
[edit services lrf profile profile-name template template-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure the template types for the template, which specify the data fields to include. You must configure at least one type, and you can configure multiple types.

If Next Gen Services is enabled, then the template types **dns**, **ifl-subscriber**, **ipv4-extended**, **ipv6-extended**, **mobile-subscriber**, **video**, and **wireline-subscriber** are not available.

Options

template-type—Template type. You must configure at least one of the following types, and you can configure multiple types:

- **device-data**—Use data fields specific to the device collecting the logging feed.
- **dns**—(Not available if Next Gen Services is enabled) Use the DNS response time data field.
- **flow-id**—Use the Flow ID data field.
- **http**—Use data fields for the HTTP metadata from header fields.
- **ifl-subscriber**—(Not available if Next Gen Services is enabled) Use data fields specific to interface-based subscribers.
- **ipflow**—Use data fields for the uplink and downlink octets and bytes.
- **ipflow-extended**—Use data fields for the service set name, routing instance, and payload timestamps.
- **ipflow-tcp**—Use data fields for TCP-related timestamps.
- **ipflow-tcp-ts**—Use IBM-specific data fields for TCP-related timestamps. When configuring a **ipflow-tcp-ts** template, configure **vendor-support ibm** at the `[edit services lrf profile profile-name]` hierarchy level to avoid a commit warning.
- **ipflow-ts**—Use data fields for the flow start and end timestamps.
- **ipv4**—Use data fields for the basic source and destination IPv4 information.

- **ipv4-extended**—(Not available if Next Gen Services is enabled) Use data fields for the elements of IPv4 extended fields.
- **ipv6**—Use data fields for the basic source and destination IPv6 information.
- **ipv6-extended**—(Not available if Next Gen Services is enabled) Use data fields for the elements of IPv6 extended fields.
- **l7-app**—Use data fields for the Layer 7 application.
- **mobile-subscriber**—(Not available if Next Gen Services is enabled) Use data fields specific to mobile subscribers.
- **pcc**—Use the PCC rule name data field.
- **status-code-dist**—Use data fields for the HTTP or DNS status codes.
- **subscriber-data**—Use data fields for Generic Subscriber information that can be included with wireless (mobile) subscribers or wireline subscribers.
- **transport-layer**—Use data fields for the transport layer.
- **video**—(Not available if Next Gen Services is enabled) Use data fields for video traffic.
- **wireline-subscriber**—(Not available if Next Gen Services is enabled) Use the UserName data field for wireline subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management | 459](#)

term (Captive Portal Content Delivery)

Syntax

```
term term-name{
  from {
    destination-address address <except>;
  }
  then {
    accept;
    insert tag-name tag-name tag-value tag-value;
    redirect url;
    rewrite destination-address address <destination-port port-number>;
    syslog;
  }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule rule-name],
[edit services captive-portal-content-delivery rule rule-name]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the **[edit dynamic-profiles *profile-name* services captive-portal-content-delivery **rule** *rule-name*]** hierarchy level added in Junos OS Release 17.2R1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

Description

Define the term match and action properties for the captive portal content delivery rule. Use the statement at the **[edit services...]** hierarchy level for static CPCD. Use the statement at the **[edit dynamic-profiles *profile-name* services...]** hierarchy level for converged services CPCD.

Options

term-name—Identifier for the term.

The remaining statement is explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

HTTP Redirect Service Overview	473
Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services	484
Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services	494
Configuring Routing Engine-Based, Static HTTP Redirect Services	505
Configuring Routing Engine-Based, Converged HTTP Redirect Services	519

term (Dynamic Profiles)

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 | 328](#)

[Configuring Terms for Fast Update Filters | 336](#)

[Fast Update Filter Match Conditions | 334](#)

[Fast Update Filter Actions and Action Modifiers | 335](#)

[Parameterized Filters Overview | 249](#)

[Firewall Filter Terminating and Nonterminating Actions for Protocol-Independent Traffic in Dynamic Service Profiles | 285](#)

[Parameterized Filter Nonterminating and Terminating Actions and Modifiers | 277](#)

[Parameterized Filter Match Conditions for IPv4 Traffic | 263](#)

[Parameterized Filter Match Conditions for IPv6 Traffic | 270](#)

then (Captive Portal Content Delivery)

Syntax

```
then {
  accept;
  insert tag-name tag-name tag-value tag-value;
  redirect url;
  rewrite destination-address address <destination-port port-number>;
  syslog;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule rule-name term term-name],
[edit services captive-portal-content-delivery rule rule-name term term-name]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the **[edit dynamic-profiles *profile-name* services captive-portal-content-delivery rule *rule-name* **term** *term-name*]** hierarchy level added in Junos OS Release 17.2R1.

insert option added in Junos OS Release 19.1R1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

Description

Define the term actions and any optional action modifiers for the captive portal content delivery rule. Use the statement at the **[edit services...]** hierarchy level for static CPCD. Use the statement at the **[edit dynamic-profiles *profile-name* services...]** hierarchy level for converged services CPCD.

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.
- **insert**—Insert a tag with a specified value that is added to the HTTP packet header when the term is matched. You must insert the following action modifiers.
 - **tag *tag-name***—Name of the tag inserted in the packet header, with a maximum size of 127 characters. The tag name is case-sensitive, such that **tag ABCD** and **tag abcd** are processed as different names:
 - **tag-value *tag-value***—Value of the specified tag, with a maximum size of 127 characters. You can specify a customer value or use one of four predefined options that causes the associated value to be inserted:

- **custom**—Value defined by user, with a maximum size of 127 characters.
- **hostname**—Hostname of the router.
- **subscriber-ip**—Subscriber's IPv4 address.
- **subscriber-ipv6**—Subscriber's IPv6 address.
- **subscriber-mac-addr**—MAC address of the subscriber.
- **redirect**—Redirect the packet and all subsequent packets in flows that match the rules. You can optionally configure the following action modifier:
 - **url**— URL destination for the redirected packet. The URL must begin with **http://** or **https://**.

rewrite—Rewrite the packet and all subsequent packets in flows that match the rules. You can optionally configure one or both of the following action modifiers:

- **destination-address address**—Destination address for the rewritten packet.
- **destination-port port-number**—(Optional) Destination port for the rewritten packet.

syslog—Log information about the packet to a system log file.

action—Actions to accept, redirect, or rewrite packets and all subsequent packets in flows that match the rules.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[HTTP Redirect Service Overview | 473](#)

Firewall Filter Match Conditions Based on Address Fields

[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

then (LRF rule)

Syntax

```
then {  
  report {  
    collector collector-name;  
    template template-name;  
    time-limit time-interval;  
    volume-limit volume;  
  }  
}
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure the actions to take if the LRF rule is matched.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers](#) | 460

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management](#) | 459

then (PCC Rules)

Syntax

```
then {  
    pcc-action-profile profile-name;  
}
```

Hierarchy Level

```
[edit unified-edge pcef pcc-rules rule-name],  
[edit services pcef pcc-rules rule-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-rules rule-name]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the policy and charging control (PCC) action profile for a PCC rule. The PCC action profile specifies the actions to apply to subscriber traffic that matches any of the **from** statements in the PCC rule. A PCC rule configuration must include the **then** statement and a PCC action profile. The referenced PCC action profile must be configured.

If you are using Junos OS Subscriber Aware, specify the name of the PCC action profile at the **[edit unified-edge pcef pcc-rules rule-name]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the name of the PCC action profile at the **[edit services pcef pcc-rules rule-name]** hierarchy level.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Policy and Charging Control Rules | 418](#)

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

[Configuring Policy and Charging Control Action Profiles for Subscriber Management | 416](#)

three-color-policer (Configuring)

Syntax

```
three-color-policer policer-name | uid {
  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.

Description

Configure a three-color policer in static firewall filters or dynamic firewall filters in a dynamic client profile or a dynamic service profile.

Options

policer-name—Name of the three-color policer. Reference this name when you apply the policer to an interface.

uid—When you configure a policer at the **[edit dynamic-profiles]** hierarchy level, you must assign a variable UID as the policer name.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Configuring and Applying Tricolor Marking Policers</i>
<i>Three-Color Policer Configuration Guidelines</i>
<i>Basic Single-Rate Three-Color Policers</i>
<i>Basic Two-Rate Three-Color Policers</i>
<i>Two-Color and Three-Color Logical Interface Policers</i>
<i>Two-Color and Three-Color Physical Interface Policers</i>
<i>Two-Color and Three-Color Policers at Layer 2</i>

time-limit (LRF Rule)

Syntax

```
time-limit time-interval;
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name then report]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure the time limit to be used for reporting. The template that the LRF rule is using must have **trigger-type time** configured.

Options

time-interval—The time limit in seconds.

Range: 60 through 1800

Default: 300

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

[Configuring Logging and Reporting for Subscriber Management | 459](#)

Configuring Logging and Reporting for Junos OS Subscriber Aware

tracoptions (Analytics Agent)

Syntax

```
tracoptions {  
    file filename;  
    flag (debug | error | info | trace);  
}
```

Hierarchy Level

[edit services analytics [agent](#)]

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Statement introduced in Junos OS Release 18.4R1 on PTX Series routers.

Description

Configure tracing operations for Network Telemetry Framework (NTF) agent. You can specify the name of the file where the NTF agent log messages are stored. You can also specify a severity level for messages to be logged. 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 **trace**. As levels become more restrictive, fewer messages are logged.

NOTE: Although the syntax uses the keyword **flag**, its function in this statement corresponds to the **level** keyword used for other **tracoptions** statements.

Options

file *filename*—Name of the file to receive the output of the tracing operation. The file is stored in the `/var/log/` directory of your device.

Default: ntf-agent

flag (debug | error | info | trace)—Specify the severity level for messages to be logged. The order of severity, from most to least severe is as follows:

error > info > debug > trace

- **debug**—Match debug messages.
- **error**—Match error messages. This is the most restrictive level.
- **info**—Match informational messages.

- **trace**—Match all messages.

Default: error

Required Privilege Level

system

RELATED DOCUMENTATION

[IPFIX Mediation on the BNG | 667](#)

Configuring NTF Agent

traceoptions (Captive Portal Content Delivery)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level

[edit services [captive-portal-content-delivery](#)]

Release Information

Statement introduced in Junos OS Release 10.4.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

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.
- **clicommand**—Trace CLI command operations.
- **configuration**—Trace home agent state machine operations.
- **general**—Trace general operations.
- **gres**—Trace graceful routing switchover operations.
- **ipc**—Trace Inter-Process Communication (IPC) messages between the PIC and the Routing Engine.
- **rtsock**—Trace routing socket operations.
- **rules**—Trace rules operations.
- **ssets**—Trace service sets 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

[HTTP Redirect Service Overview | 473](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services | 484](#)

[Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services | 494](#)

[Configuring Routing Engine-Based, Static HTTP Redirect Services | 505](#)

[Configuring Routing Engine-Based, Converged HTTP Redirect Services | 519](#)

traceoptions (TCP Port Forwarding)

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);
}
```

Hierarchy Level

[edit system processes [tcp-forwarding](#)]

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Define tracing operations for TCP port forwarding 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 events
- **configuration**—Trace configuration events
- **connection**—Trace TCP connection events
- **init**—Trace TCP port forwarding initialization events

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 all levels.
- **error**—Match error messages.

- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

Default: error

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—(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: *sizek* to specify KB, *sizem* to specify MB, or *sizeg* 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 TCP Port Forwarding Events for Troubleshooting | 662](#)

[TCP Port Forwarding for Remote Device Management | 655](#)

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>;
  max-burst-size cells;
  overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
  peak-rate rate;
  scheduler-map map-name;
  shaping-rate (percent percentage | rate | predefined-variable) <burst-size bytes>;
  shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
  shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
  shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
  shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
  shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
  shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
  shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
  shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
  shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
  sustained-rate rate;
}
```

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 for use in a dynamic client profile or a dynamic service profile.

Options

profile-name—Name of the traffic-control profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Guidelines for Configuring Dynamic CoS for Subscriber Access | 40](#)

[Configuring Traffic Scheduling and Shaping for Subscriber Access | 47](#)

[Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers | 50](#)

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 6,400,000,000,000 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 | 40](#)

[Configuring Schedulers in a Dynamic Profile for Subscriber Access | 52](#)

[scheduler | 1134](#)

trigger-type (LRF Profile)

Syntax

```
trigger-type (session-close | volume);
```

Hierarchy Level

```
[edit services lrf profile profile-name template template-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure the type of trigger that causes the generation of data records and transmission to the collector. You can only configure one type of trigger.

Default

If you do not include the **trigger-type** statement, the default trigger is **session-close**.

Options

session-close—Use the closing of the data session to cause the generation of data records and transmission to the collector.

volume—Use a data volume limit to cause the generation of data records and transmission to the collector. The data volume limit value is configured in the LRF rule.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

[Configuring Logging and Reporting for Junos OS Subscriber Aware](#)

[Configuring Logging and Reporting for Subscriber Management | 459](#)

tunnel-services (Chassis)

Syntax

```
tunnel-services {  
    bandwidth bandwidth-value;  
    tunnel-only;  
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number]
```

Release Information

Statement introduced in Junos OS Release 8.2.

Statement introduced in Junos OS Release 12.3X54 for ACX Series routers.

Description

For MX Series 5G Universal Routing Platforms, configure the amount of bandwidth for tunnel services.

For ACX Series routers, configure the amount of bandwidth for tunnel services. Only bandwidths of 1 Gbps and 10 Gbps are supported for ACX routers.

For M7i, M10i, M120, M320, T Series and TX Matrix routers with IQ2 PICs and IQ2E PICs, configure support for per unit scheduling for GRE tunnels. You can specify the IQ2 and IQ2E PICs to work exclusively in tunnel mode or as a regular PIC. The default setting uses IQ2 and IQ2E PICs as a regular PIC. If you do not configure the **tunnel-only** option, the IQ2 and IQ2E PICs operate as regular PICs. For M7i, M10i, M120, M320, T Series and TX Matrix routers with IQ2 PICs and IQ2E PICs, you can use the **tunnel-only** option to specify that an IQ2 or IQ2E PIC work in tunnel mode only.

NOTE: Bandwidth rates of 20 gigabits per second and 40 gigabits per second require use of an MX Series router with the 100-Gigabit Ethernet Modular Port Concentrator (MPC) and the 100-Gigabit CFP MIC.

NOTE: On MX80 routers and MX Series routers with Trio-based FPCs, when ingress queuing is enabled for a PIC, tunnel services and inline services are not supported on the same PIC.

Options

tunnel-only (Optional)—For M7i, M10i, M120, M320, T Series and TX Matrix routers with IQ2 PICs and IQ2E PICs, specify that an IQ2 or IQ2E PIC work in tunnel mode only.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Tunnel Interfaces on a Gigabit Ethernet 40-Port DPC

Example: Configuring Tunnel Interfaces on a 10-Gigabit Ethernet 4-Port DPC

Example: Configuring Tunnel Interfaces on the MPC3E

[bandwidth \(Tunnel Services\)](#) | 743

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.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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](#) | [784](#)

[color-blind](#) | [786](#)

[single-rate](#) | [1162](#)

type (Application Identification)

Syntax

```
type type;
```

Hierarchy Level

```
[edit services application-identification application application-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Specify the type of application, such as FTP or HTTP.

Options

type—Type of application such as FTP or HTTP.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview](#) | [431](#)

[Configuring Custom Application Signatures](#) | [435](#)

type (ICMP Mapping for Application Identification)

Syntax

```
type icmp-type;
```

Hierarchy Level

```
[edit services application-identification application application-name icmp-mapping]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Match an ICMP type value to create a custom application signature.

Options

value—ICMP code value.

Range: 0 through 254

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

uid (Dynamic Profiles)

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](#)

uid-reference

Syntax

```
uid-reference;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name variables variable-name]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Description

When you configure a unique ID (UID) variable, include this statement to specify that the value for the UID is supplied by RADIUS 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

[Unique Identifiers for Firewall Variables | 250](#)

[Configuring Unique Identifiers for Parameterized Filters | 253](#)

[Dynamic Variables Overview](#)

unit (Dynamic Profiles Standard Interface)

Syntax

```

unit logical-unit-number {
    actual-transit-statistics;
    auto-configure {
        agent-circuit-identifier {
            dynamic-profile profile-name;
        }
        line-identity {
            include {
                accept-no-ids;
                circuit-id;
                remote-id;
            }
            dynamic-profile profile-name;
        }
    }
    dial-options {
        ipsec-interface-id name;
        l2tp-interface-id name;
        (shared | dedicated);
    }
    encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux | atm-mlppp-llc | atm-nlpid
        | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc |
        ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | frame-relay-ccc | frame-relay-ppp
        | frame-relay-tcc | frame-relay-ether-type | frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end |
        multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc |
        vlan-vpls);
    family family {
        address address;
        demux-destination,
        filter {
            adf {
                counter;
                input-precedence precedence;
                not-mandatory;
                output-precedence precedence;
                rule rule-value;
            }
            input filter-name {
                precedence precedence;
                shared-name filter-shared-name;
            }
        }
    }
}

```

```

output filter-name {
    precedence precedence;
    shared-name filter-shared-name;
}
}
max-sessions number;
max-sessions-vsa-ignore;
rpf-check {
    fail-filter filter-name;
    mode loose;
}
service {
    input {
        service-set service-set-name {
            service-filter filter-name;
        }
        post-service-filter filter-name;
    }
    input-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (push | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
    output {
        service-set service-set-name {
            service-filter filter-name;
        }
    }
    output-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (pop | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
}
service-name-table table-name
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>;
unnumbered-address interface-name <preferred-source-address address>;
}

```

```

filter {
  input filter-name {
    shared-name filter-shared-name;
  }
  output filter-name {
    shared-name filter-shared-name;
  }
}
host-prefix-only;
keepalives {
  interval seconds;
}
ppp-options {
  aaa-options aaa-options-name;
  authentication [ authentication-protocols ];
  chap {
    challenge-length minimum minimum-length maximum maximum-length;
    local-name name;
  }
  ignore-magic-number-mismatch;
  initiate-ncp (dual-stack-passive | ipv6 | ip)
  ipcp-suggest-dns-option;
  mru size;
  mtu (size | use-lower-layer);
  on-demand-ip-address;
  pap;
  peer-ip-address-optional;
  local-authentication {
    password password;
    username-include {
      circuit-id;
      delimiter character;
      domain-name name;
      mac-address;
      remote-id;
    }
  }
}
service {
  pcef pcef-profile-name {
    activate rule-name | activate-all;
  }
}

```

```
targeted-options {
  backup backup;
  group group;
  primary primary;
  weight ($junos-interface-target-weight | weight-value);
}
vlan-id number;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
```

Hierarchy Level

[edit [dynamic-profiles](#) *profile-name* [interfaces](#) *interface-name*]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—The specific unit number of the interface you want to assign to the dynamic profile, or one of the following predefined variables:

- **\$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **\$junos-interface-unit**—The unit number variable on a dynamic underlying VLAN interface for which you want to enable the creation of dynamic VLAN subscriber interfaces based on the ACI.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information

Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information

Agent Circuit Identifier-Based Dynamic VLANs Overview

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);
  report-ingress-shaping-rate bps;
  rewrite-rules {
    dscp (rewrite-name | default);
    dscp-ipv6 (rewrite-name | default);
    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
    inet-precedence (rewrite-name | default);
  }
}
}
```

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 | 40](#)

[Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile | 217](#)

[Configuring an Interface Set of Subscribers in a Dynamic Profile | 192](#)

url

Syntax

```
url url-name;
```

Hierarchy Level

```
[edit unified-edge pcef pcc-action-profiles profile-name redirect],  
[edit services pcef pcc-action-profiles profile-name redirect]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support at the **[edit services pcef pcc-action-profiles *profile-name* redirect]** hierarchy level introduced for Junos OS Broadband Subscriber Management in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services for Junos OS Broadband Subscriber Management introduced in Junos OS Release 19.3R2 on MX Series.

Description

Specify the URL name that you want a PCC action profile to use for performing HTTP redirection. If you configure this, the PCC action profile can only be used in PCC rules that match only HTTP-based applications and all flows.

If you are using Junos OS Subscriber Aware, specify the URL name at the **[edit unified-edge pcef pcc-action-profiles *profile-name* redirect]** hierarchy level.

If you are using Junos OS Broadband Subscriber Management, specify the URL name at the **[edit services pcef pcc-action-profiles *profile-name* redirect]** hierarchy level.

Options

url-name—URL for the HTTP redirect.

Required Privilege Level

For Junos OS Subscriber Aware:

unified-edge—To view this statement in the configuration.

unified-edge-control—To add this statement to the configuration.

For Junos OS Broadband Subscriber Management:

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Policy and Charging Control Action Profiles For Junos OS Subscriber Aware

[Configuring Policy and Charging Control Action Profiles for Subscriber Management](#) | 416

user (Access)

Syntax

```

user user-name {
  authentication {
    encrypted-password encrypted-password;
    no-public-keys;
    ssh-eccdsa name {
      from host-list;
    }
    ssh-ed25519 name {
      from host-list;
    }
    ssh-rsa name {
      from host-list;
    }
  }
  class class-name;
  cli {
    prompt prompt;
  }
  full-name complete-name;
  uid uid;
}

```

Hierarchy Level

[edit system [login](#)]

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Statement **no-public-keys** introduced in Junos OS Release 15.1.

Statement **cli** introduced in Junos OS 17.3.

Description

Configure access permission for individual users. Starting in Junos OS Release 18.3, the **ssh-dsa** hostkey algorithm is deprecated— rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

Options

authentication—Specify one or more authentication methods that a user can use to log in to the router or switch. You can assign multiple authentication methods to a single user.

encrypted-password— Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.



CAUTION: Do not use the **encrypted-password** option unless the password is *already* encrypted, and you are entering the encrypted version of the password.

If you accidentally configure the **encrypted-password** statement with a plain-text password or with blank quotation marks (" "), you will not be able to log in to the device as this user.

Range: You cannot configure a blank password for **encrypted-password** using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

no-public-keys— Disables ssh public key authentication for the user specified. If the no-public-keys statement is specified at the [edit system services ssh] hierarchy level, public key authentication is disabled for all users on the device.

ssh-ecdsa public-key— SSH version 2 authentication. Specify the ECDSA public key. You can specify one or more public keys for each user.

from host-list— Specify a pattern-list of allowed hosts.

ssh-ed25519 public-key— SSH version 2 authentication. Specify the ED25519 public key. You can specify one or more public keys for each user.

from host-list— Specify a pattern-list of allowed hosts.

ssh-rsa public-key— SSH version 2 authentication. Specify the RSA public key. You can specify one or more public keys for each user.

from host-list— Specify a pattern-list of allowed hosts.

class class-name— Assign a user to a login class. You must assign each user to a login class. Specify one of the classes defined at the [edit system login class] hierarchy level.

cli— Set the CLI prompt specified for a specified login user or specified login class. The prompt set for the login user has precedence.

prompt prompt— Specify the prompt string you want to see displayed in the CLI prompt.

full-name complete-name— Specify the user's complete name. If the name contains spaces, enclose it in quotation marks. Do not include colons or commas.

uid *uid-value*— Numeric identifier associated with the user account, either assigned by an administrator or assigned automatically when you commit the user configuration. It is used by applications that request numeric identifiers, such as some RADIUS queries, or secure applications, such as flow-tap monitoring. This value must be unique on the router or switch.

Default: If you do not assign a UID to a user, the software assigns one when you commit the configuration, preferring the lowest available number.

Range: 100 through 64000

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

	<i>Configuring Junos OS User Accounts by Using a Configuration Group</i>
	<i>set cli prompt</i>
	class 767
	<i>root-authentication</i>

vendor-support

Syntax

```
vendor-support ibm;
```

Hierarchy Level

```
[edit services lrf profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 17.2 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure support for any vendor-specific template types. Currently, the only vendor-specific template type is **ipflow-tcp-ts**, for which you configure **vendor-specific ibm**.

If you do not configure **vendor-specific ibm**, a warning appears when you commit the configuration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management | 459](#)

[Logging and Reporting Function for Subscribers | 441](#)

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

[Dynamic IGMP Configuration Overview | 393](#)

[Configuring Dynamic DHCP Client Access to a Multicast Network | 395](#)

[Changing the IGMP Version](#)

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

[Dynamic MLD Configuration Overview](#) | 401

Modifying the MLD Version

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 | 40](#)

[Applying a Classifier to a Subscriber Interface in a Dynamic Profile | 221](#)

classifiers (Definition)

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](#) | 40

[Applying a Rewrite Rule Definition to a Subscriber Interface in a Dynamic Profile](#) | 219

[rewrite-rules](#)

volume-limit (LRF Rule)

Syntax

```
volume-limit volume;
```

Hierarchy Level

```
[edit services lrf profile profile-name rule lrf-rule-name then report]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support for Next Gen Services introduced in Junos OS Release 19.3R1 on MX Series.

Description

Configure the data volume limit to be used for reporting. The template that the LRF rule is using must have **trigger-type volume** configured.

Options

volume—Data volume, in megabytes.

Range: 1 through 1024

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LRF Profile for Subscribers | 460](#)

Configuring Logging and Reporting for Junos OS Subscriber Aware

[Configuring Logging and Reporting for Subscriber Management | 459](#)

Operational Commands

IN THIS CHAPTER

- [clear firewall | 1256](#)
- [clear igmp membership | 1259](#)
- [clear interfaces statistics | 1263](#)
- [clear mld membership | 1265](#)
- [clear remote-device-management statistics | 1267](#)
- [clear services application-identification application-system-cache | 1269](#)
- [clear services application-identification statistics | 1270](#)
- [clear services captive-portal-content-delivery statistics | 1273](#)
- [clear services lrf collector statistics | 1275](#)
- [clear services lrf statistics | 1276](#)
- [clear tcp-forwarding connections | 1277](#)
- [clear tcp-forwarding statistics | 1280](#)
- [request interface rebalance \(Aggregated Ethernet for Subscriber Management\) | 1284](#)
- [request network-access aaa subscriber add session-id | 1285](#)
- [request network-access aaa subscriber delete session-id | 1287](#)
- [request network-access aaa subscriber modify session-id | 1290](#)
- [request network-access aaa subscriber set session-id | 1292](#)
- [request services application-identification application | 1294](#)
- [request services application-identification download | 1296](#)
- [request services application-identification download status | 1297](#)
- [request services application-identification group | 1298](#)
- [request services application-identification install | 1300](#)
- [request services application-identification install status | 1302](#)
- [request services application-identification proto-bundle-status | 1303](#)
- [request services application-identification uninstall | 1304](#)
- [request services application-identification uninstall status | 1305](#)
- [request services remote-device-management reconfigure service-device | 1306](#)
- [request services remote-device-management reload-dictionary | 1308](#)

- [show class-of-service | 1310](#)
- [show class-of-service adjustment-control-profile | 1313](#)
- [show class-of-service interface | 1315](#)
- [show class-of-service interface-set | 1355](#)
- [show class-of-service scheduler-hierarchy interface | 1358](#)
- [show class-of-service scheduler-hierarchy interface-set | 1361](#)
- [show class-of-service scheduler-map | 1363](#)
- [show class-of-service traffic-control-profile | 1367](#)
- [show dynamic-profile session | 1372](#)
- [show firewall | 1378](#)
- [show firewall log | 1389](#)
- [show firewall templates-in-use | 1393](#)
- [show igmp group | 1395](#)
- [show igmp interface | 1400](#)
- [show interfaces statistics | 1405](#)
- [show interfaces targeting \(Aggregated Ethernet for Subscriber Management\) | 1422](#)
- [show mld group | 1424](#)
- [show mld interface | 1429](#)
- [show network-access aaa subscribers session-id | 1434](#)
- [show services analytics agent | 1442](#)
- [show remote-device-management service-devices | 1445](#)
- [show remote-device-management services | 1453](#)
- [show remote-device-management statistics | 1456](#)
- [show remote-device-management subscribers | 1461](#)
- [show remote-device-management summary | 1465](#)
- [show services application-identification application | 1469](#)
- [show services application-identification application-system-cache | 1477](#)
- [show services application-identification commit-status \(Next Gen Services\) | 1482](#)
- [show services application-identification counter | 1484](#)
- [show services application-identification group | 1487](#)
- [show services application-identification statistics application-groups | 1492](#)
- [show services application-identification statistics applications | 1494](#)
- [show services application-identification status | 1496](#)
- [show services application-identification version | 1499](#)

- [show services captive-portal-content-delivery](#) | 1500
- [show services lrf collector statistics](#) | 1506
- [show services lrf rule statistics](#) | 1508
- [show services lrf statistics](#) | 1510
- [show services lrf template](#) | 1512
- [show services pcef pic](#) | 1515
- [show services pcef subscribers](#) | 1517
- [show services service-sets summary](#) | 1525
- [show subscribers](#) | 1527
- [show subscribers summary](#) | 1576
- [show tcp-forwarding status](#) | 1585

clear firewall

List of Syntax

[Syntax on page 1256](#)

[Syntax \(EX Series Switches\) on page 1256](#)

Syntax

```
clear firewall (all | counter counter-name | filter filter-name | log (all | logical-system-name) | logical-system
logical-system-name)
```

Syntax (EX Series Switches)

```
clear firewall (all | counter counter-name | filter filter-name | log (all | logical-system-name) | policer counter (all |
counter-id counter-index))
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

logical-system option introduced in Junos OS Release 9.3.

log option introduced before Junos OS Release 11.4.

Description

Clear statistics about configured firewall filters.

When you clear the counters of a filter, this impacts not only the counters shown by the CLI, but also the ones tracked by SNMP2.

Subscriber management uses firewall filters to capture and report the volume-based service accounting counters that are used for subscriber billing. The **clear firewall** command also clears the service accounting counters that are reported to the RADIUS accounting server. For this reason, you must be cautious in specifying which firewall statistics you want to clear.

NOTE: The **clear firewall** command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover (GRES).

If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the **prefix-action** action on matched packets, wait at least 5 seconds before you enter the **show firewall prefix-action-stats** command. A 5-second pause between issuing the **clear firewall** and **show firewall prefix-action-stats** commands avoids a possible timeout of the **show firewall prefix-action-stats** command.

Options

all—Clear the packet and byte counts for all filters. On EX Series switches, this option also clears the packet counts for all policer counters.

counter *counter-name*—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.

filter *filter-name*—Clear the packet and byte counts for the specified firewall filter.

log (all | *logical-system-name*)—Clear log entries for IPv4 firewall filters that have **then log** as an action. Use **log all** to clear all log entries or **log *logical-system-name*** to clear log entries for the specified logical system.

logical-system *logical-system-name*—Clear the packet and byte counts for the specified logical system.

policer counter (all | counter-id *counter-index*)—(EX8200 switches only) Clear all policer counters using the **policer counter all** command, or clear a specific policer counter using the **policer counter counter-id *counter-index*** command. The value of *counter-index* can be 0, 1, or 2.

Required Privilege Level

clear

RELATED DOCUMENTATION

[show firewall](#) | [1378](#)

List of Sample Output

[clear firewall all on page 1257](#)

[clear firewall \(counter counter-name\) on page 1257](#)

[clear firewall \(filter filter-name\) on page 1258](#)

[clear firewall \(policer counter all\) \(EX8200 Switch\) on page 1258](#)

[clear firewall \(policer counter counter-id counter-index\) \(EX8200 Switch\) on page 1258](#)

Sample Output

clear firewall all

```
user@host> clear firewall all
```

clear firewall (counter counter-name)

```
user@host> clear firewall counter port-filter-counter
```

clear firewall (filter filter-name)

```
user@host> clear firewall filter ingress-port-filter
```

clear firewall (policer counter all) (EX8200 Switch)

```
user@switch> clear firewall policer counter all
```

clear firewall (policer counter counter-id counter-index) (EX8200 Switch)

```
user@switch> clear firewall policer counter counter-id 0
```


clear igmp membership

List of Syntax

[Syntax on page 1259](#)

[Syntax \(EX Series Switch and the QFX Series\) on page 1259](#)

Syntax

```
clear igmp membership
<all>
<group address-range>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
clear igmp membership
<group address-range>
<interface interface-name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.3 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear Internet Group Management Protocol (IGMP) group members.

Options

all—Clear IGMP members for groups and interfaces in the master instance.

group *address-range*—(Optional) Clear all IGMP members that are in a particular address range. An example of a range is **233.252/16**. If you omit the destination prefix length, the default is **/32**.

interface *interface-name*—(Optional) Clear all IGMP group members on an interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

RELATED DOCUMENTATION

show igmp group 1395
show igmp interface 1400

List of Sample Output

- [clear igmp membership all on page 1260](#)
- [clear igmp membership interface on page 1261](#)
- [clear igmp membership group on page 1262](#)

Output Fields

See [show igmp group](#) for an explanation of output fields.

Sample Output

clear igmp membership all

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

user@host> **show igmp group**

Interface	Group	Last Reported	Timeout
so-0/0/0	198.51.100.253	203.0.113.1	186
so-0/0/0	198.51.100.254	203.0.113.1	186
so-0/0/0	198.51.100.255	203.0.113.1	187
so-0/0/0	198.51.100.240	203.0.113.1	188
local	198.51.100.6	(null)	0
local	198.51.100.5	(null)	0
local	198.51.100.25	(null)	0
local	198.51.100.22	(null)	0
local	198.51.100.2	(null)	0
local	198.51.100.13	(null)	0

user@host> **clear igmp membership all**

Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0

user@host> **show igmp group**

Interface	Group	Last Reported	Timeout
local	198.51.100.6	(null)	0
local	198.51.100.5	(null)	0
local	198.51.100.254	(null)	0
local	198.51.100.255	(null)	0
local	198.51.100.2	(null)	0
local	198.51.100.13	(null)	0

clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
so-0/0/0	198.51.100.253	203.0.113.1	210
so-0/0/0	198.51.100.200	203.0.113.1	210
so-0/0/0	198.51.100.255	203.0.113.1	215
so-0/0/0	198.51.100.254	203.0.113.1	216
local	198.51.100.6	(null)	0
local	198.51.100.5	(null)	0
local	198.51.100.254	(null)	0
local	198.51.100.255	(null)	0
local	198.51.100.2	(null)	0
local	198.51.100.13	(null)	0

```
user@host> clear igmp membership interface so-0/0/0
```

```
Clearing Group Membership Info for so-0/0/0
```

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
local	198.51.100.6	(null)	0
local	198.51.100.5	(null)	0
local	198.51.100.254	(null)	0
local	198.51.100.255	(null)	0
local	198.51.100.2	(null)	0
local	198.51.100.13	(null)	0

clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
so-0/0/0	198.51.100.253	203.0.113.1	210
so-0/0/0	198.51.100.25	203.0.113.1	210
so-0/0/0	198.51.100.255	203.0.113.1	215
so-0/0/0	198.51.100.254	203.0.113.1	216
local	198.51.100.6	(null)	0
local	198.51.100.5	(null)	0
local	198.51.100.254	(null)	0
local	198.51.100.25	(null)	0
local	198.51.100.2	(null)	0
local	198.51.100.13	(null)	0

```
user@host> clear igmp membership group 233.252/16
```

```
Clearing Group Membership Range 198.51.100.0/16 on so-0/0/0
Clearing Group Membership Range 198.51.100.0/16 on so-1/0/0
Clearing Group Membership Range 198.51.100.0/16 on so-2/0/0
```

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
so-0/0/0	198.51.100.255	203.0.113.1	231
so-0/0/0	198.51.100.254	203.0.113.1	233
so-0/0/0	198.51.100.253	203.0.113.1	236
local	198.51.100.6	(null)	0
local	198.51.100.5	(null)	0
local	198.51.100.254	(null)	0
local	198.51.100.255	(null)	0
local	198.51.100.2	(null)	0
local	198.51.100.13	(null)	0

clear interfaces statistics

Syntax

```
clear interfaces statistics (all | interface-name)
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 19.2R1 for QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.

Description

Set interface statistics to zero. If you issue the **clear interfaces statistics *interface-name*** command and then perform a graceful Routing Engine switchover, the interface statistics are not cleared on the new master. Reissue the command to clear the interface statistics again.

Starting in Junos OS Release 17.3R1, this command supports the clearing of Packet Forwarding Engine accounting statistics on logical interfaces configured with accounting options. On these interfaces, the current statistics values are stored as the new current baseline values and then the counters are reset to zero. If the **allow-clear** statement is included in the interface profile, then the cleared statistics values are reported to the accounting options flat file associated with the interface. Reporting is disabled by default; if **allow-clear** is not configured, then the CLI displays cleared statistics counters, but they are not reported to the flat file.

Starting in Junos OS Release 19.1R1, this command supports the clearing of unicast Reverse Path Forwarding (RPF) statistics.

Options

all—Set statistics on all interfaces to zero.

interface-name—Set statistics on a particular interface to zero.

Required Privilege Level

clear

List of Sample Output

[clear interfaces statistics on page 1264](#)

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

```
clear interfaces statistics
```

```
user@host> clear interfaces statistics
```

clear mld membership

Syntax

```
clear mld membership  
<all>  
<group group-name>  
<interface interface-name>  
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

Clear Multicast Listener Discovery (MLD) group membership.

Options

all—Clear MLD memberships for groups and interfaces in the master instance.

group *group-name*—(Optional) Clear MLD membership for the specified group.

interface *interface-name*—(Optional) Clear MLD group membership for the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

| [show mld group](#) | [1424](#)

List of Sample Output

[clear mld membership all on page 1266](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
clear mld membership all
```

```
user@host> clear mld membership all
```


clear remote-device-management statistics

Syntax

```
clear remote-device-management statistics  
(summary | service-devices device-name)
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Clear service statistics for all remote devices globally or statistics for a specific remote service device.

Options

service-devices *device-name*—(Optional) Clear statistics for the specified service device.

summary—(Optional) Clear service statistics for all remote devices.

Required Privilege Level

view

RELATED DOCUMENTATION

| [show remote-device-management statistics](#) | 1456

List of Sample Output

[clear remote-device-management statistics \(Service Device\) on page 1267](#)

[clear remote-device-management statistics \(Global\) on page 1268](#)

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the [show remote-device-management statistics](#) command before and after clearing the service statistics to verify the clear operation.

Sample Output

clear remote-device-management statistics (Service Device)

```
user@host> clear remote-device-management statistics service-device olt-xyz
```

clear remote-device-management statistics (Global)

`user@host> clear remote-device-management statistics summary`

clear services application-identification application-system-cache

Syntax

```
clear services application-identification application-system-cache
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Clear entries from the application system cache.

Options

This command has no options.

Required Privilege Level

clear

RELATED DOCUMENTATION

[show services application-identification application-system-cache](#) | 1477

List of Sample Output

[clear services application-identification application-system-cache on page 1269](#)

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear services application-identification application-system-cache

```
user@host> clear services application-identification application-system-cache
```

clear services application-identification statistics

Syntax

```
clear services application-identification statistics
<cumulative>
<interval>
<logical-system (logical-system-name | all | root-logical-system)>
<tenant (tenant-name | all)>
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

logical-system option introduced in Junos OS Release 18.3R1 on SRX Series.

tenant option introduced in Junos OS Release 19.4R1 on SRX Series.

Description

Clears all Junos OS application statistics such as cumulative, interval, applications, and application groups.

Options

cumulative—(Optional) Clears the cumulative application statistics.

interval—(Optional) Clears the application interval statistics. Interval statistics are displayed in Top-N format, such that the first application group displayed has the largest byte count. If this parameter is not specified, then the default is 1, which is the current interval.

logical-system *logical-system-name*—(Optional) Clears application identification statistics of the specified logical system.

logical-system *all*—(Optional) Clears application identification statistics of all the logical systems.

root-logical-system—(Optional) Clears application identification statistics of the root logical system.

tenant *tenant-name*—(Optional) Clears application identification statistics of the specified tenant system.

tenant *all*—(Optional) Clears application identification statistics of all the tenant systems.

Required Privilege Level

clear

RELATED DOCUMENTATION

[show services application-identification statistics applications](#) | 1494

[show services application-identification statistics application-groups](#) | 1492

List of Sample Output

[clear services application-identification statistics on page 1271](#)

[clear services application-identification statistics logical-system all on page 1271](#)

[clear services application-identification statistics cumulative tenant TSYS1 on page 1271](#)

[clear services application-identification statistics cumulative tenant all on page 1271](#)

[clear services application-identification statistics cumulative on page 1271](#)

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear services application-identification statistics

```
user@host> clear services application-identification statistics
```

```
appid statistics cleared
```

clear services application-identification statistics logical-system all

```
user@host> clear services application-identification statistics logical-system all
```

```
appid statistics cleared
```

clear services application-identification statistics cumulative tenant TSYS1

```
user@host> clear services application-identification statistics cumulative tenant TSYS1
```

```
appid statistics cleared
```

clear services application-identification statistics cumulative tenant all

```
user@host> clear services application-identification statistics cumulative tenant all
```

```
appid statistics cleared
```

clear services application-identification statistics cumulative

```
user@host:TSYS1> clear services application-identification statistics cumulative
```

appid statistics cleared

clear services captive-portal-content-delivery statistics

Syntax

```
clear services captive-portal-content-delivery statistics
<interface pic-name>
```

Release Information

Command introduced in Junos OS Release 10.4.

Description

Clear captive portal content delivery statistics.

Options

interface—Clear statistics by PIC name.

Required Privilege Level

clear

RELATED DOCUMENTATION

| [show services captive-portal-content-delivery](#) | 1500

Output Fields

When you enter this command, you receive feedback on the status of your request.

clear services captive-portal-content-delivery statistics

```
user@host> clear services captive-portal-content-delivery statistics interface ms-5/0/0
```

```
user@host> show services captive-portal-content-delivery statistics interface ms-5/0/0
```

```
service-set interface: ms-5/0/0

Packets received   Packets altered
0                  0
```

Note that the stats are cleared.

clear services lrf collector statistics

Syntax

```
clear services lrf collector statistics  
<collector-name>
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Clear all the LRF statistics for the specified collector. If a collector is not specified, statistics are cleared for all collectors.

Options

none—Clear LRF statistics for all collectors.

collector-name—(Optional) Clear LRF statistics for the specified collector.

Required Privilege Level

clear

RELATED DOCUMENTATION

[show services lrf collector statistics](#) | 1506

Output Fields

A message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

```
clear services lrf collector statistics
```

```
user@host> clear services lrf collector statistics coll1
```

```
Interface: ms-0/1/0, Status: LRF collector statistics successfully cleared
```

clear services lrf statistics

Syntax

```
clear services lrf statistics
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Clear all the LRF statistics.

Required Privilege Level

clear

RELATED DOCUMENTATION

[show services lrf statistics](#) | [1510](#)

Output Fields

A message is displayed on successful execution of this command; otherwise an error message is displayed.

Sample Output

```
clear services lrf statistics
```

```
user@host> clear services lrf statistics
```

```
Interface: ms-0/1/0, Status: LRF statistics successfully cleared
```

clear tcp-forwarding connections

Syntax

```
clear tcp-forwarding connections
listening-port listening-port-number listening-address ipv4-listening-address
source-address source-ipv4-address source-port source-port-number
routing-instance routing-instance-name
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Clear all TCP connections, all connections associated with a specific listening port/listening address combination, or a single connection pair represented by a specific source address/source port combination. For either combination, you can optionally specify a routing instance. If you do not specify a routing instance, the default routing instance is assumed. This command enables you to disconnect TCP port forwarding connections that are not behaving properly.

Options

listening-address *ipv4-listening-address*—IPv4 address that is part of a listening port/listening address combination. The listening address is one on the BNG that external management systems or remote devices must use when attempting to trigger connections on the listening port. You must also specify a listening port.

listening-port *port-number*—Port number that is part of a listening port/listening address combination. The listening port is one that the BNG monitors for connections to be triggered by external management systems or remote devices.

Range: 8000 through 8031

source-address *source-ipv4-address*—Source address of the triggering entity—the remote device or external management system—that appear in the TCP header.

source-port *source-port-number*—Source port of the triggering entity—the remote device or external management system—that appear in the TCP header.

Range: 1 through 65,535

routing-instance *routing-instance-name*—Name of the routing instance for the TCP mapping.

Required Privilege Level

view

RELATED DOCUMENTATION

List of Sample Output

[clear tcp-forwarding connections on page 1278](#)

[clear tcp-forwarding connections \(Listening Port and Address\) on page 1278](#)

[clear tcp-forwarding connections \(Source Address and Port\) on page 1279](#)

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear tcp-forwarding connections

```
user@host> clear tcp-forwarding connections
```

clear tcp-forwarding connections (Listening Port and Address)

The following sample output displays the TCP connection status for a specific listening port/address combination before and after the connection is cleared.

```
user@host> show tcp-forwarding status listening-port 203.0.113.50 listening-address 8002
```

```
Listening on: [default:]203.0.113.50:8002
  Status: listening
  Total Bytes Rx: 1230 Tx: 482
Forwarding to: [default:]192.0.0.4:830
  Total Bytes Rx: 482 Tx: 1230
Allowed Source Prefixes:
  198.51.100.3/32
Connections Max: 1 Active: 1
Source: 198.51.100.3:55002
  Listening: connected Bytes Rx: 0 Tx: 0
  Forwarding: connected Bytes Rx: 0 Tx: 0
```

```
user@host> clear tcp-forwarding connections listening-port 8002 listening-address 203.0.113.50
```

```
user@host> show tcp-forwarding status listening-port 8002 listening-address 203.0.113.50
```

clear tcp-forwarding connections (Source Address and Port)

The following sample output displays the TCP connection status before and after the connection is cleared.

```
user@host> show tcp-forwarding status
```

```
Listening on: [default:]203.0.113.50:8020
  Status: listening
  Total Bytes Rx: 292 Tx: 112
Forwarding to: [default:]198.51.100.1:49
  Total Bytes Rx: 112 Tx: 292
Allowed Source Prefixes:
  192.0.0.1/24
Connections Max: 4 Active: 3
  Source: 192.0.0.2:55000
    Listening: connected Bytes Rx: 380 Tx: 223
    Forwarding: connected Bytes Rx: 223 Tx: 380
Source: 192.0.0.3:55000
  Listening: connected Bytes Rx: 855 Tx: 411
  Forwarding: connected Bytes Rx: 411 Tx: 855
Source: 192.0.0.4:56022
  Listening: connected Bytes 642 Tx: 350
  Forwarding: connected Bytes Rx: 350 Tx: 642
```

```
user@host> clear tcp-forwarding connections source-address 192.0.0.2 source-port 55000
```

```
user@host> show tcp-forwarding status
```

```
Listening on: [default:]203.0.113.50:8020
  Status: listening
  Total Bytes Rx: 292 Tx: 112
Forwarding to: [default:]198.51.100.1:49
  Total Bytes Rx: 112 Tx: 292
Allowed Source Prefixes:
  192.0.0.1/24
Connections Max: 4 Active: 2
Source: 192.0.0.3:55000
  Listening: connected Bytes Rx: 855 Tx: 411
  Forwarding: connected Bytes Rx: 411 Tx: 855
Source: 192.0.0.4:56022
  Listening: connected Bytes 642 Tx: 350
  Forwarding: connected Bytes Rx: 350 Tx: 642
```

clear tcp-forwarding statistics

Syntax

```
clear tcp-forwarding statistics
listening-port listening-port-number listening-address ipv4-listening-address
source-address source-ipv4-address source-port source-port-number
routing-instance routing-instance-name
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Clear the statistics displayed by the **show tcp-forwarding status** command. You can clear statistic for all TCP mappings, for all connections associated with a specific listening port/listening address combination, or for only a single connection pair represented by a specific source address/source port combination. For either combination, you can optionally specify a routing instance. If you do not specify a routing instance, the default routing instance is assumed.

Options

listening-address *ipv4-listening-address*—IPv4 address that is part of a listening port/listening address combination. The listening address is one on the BNG that external management systems or remote devices must use when attempting to trigger connections on the listening port. You must also specify a listening port.

listening-port *port-number*—Port number that is part of a listening port/listening address combination. The listening port is one that the BNG monitors for connections to be triggered by external management systems or remote devices.

Range: 8000 through 8031

source-address *source-ipv4-address*—Source address of the triggering entity—the remote device or external management system—that appear in the TCP header.

source-port *source-port-number*—Source port of the triggering entity—the remote device or external management system—that appear in the TCP header.

Range: 1 through 65,535

routing-instance *routing-instance-name*—Name of the routing instance for the TCP mapping.

Required Privilege Level

view

RELATED DOCUMENTATION

List of Sample Output

[clear tcp-forwarding statistics \(Listening Port and Address\) on page 1281](#)

[clear tcp-forwarding statistics \(Source Address and Port\) on page 1282](#)

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output**clear tcp-forwarding statistics (Listening Port and Address)**

The following sample output displays the TCP connection status for a specific listening port/address combination before and after the connection is cleared.

user@host> **show tcp-forwarding status listening-port 8002 listening-address 203.0.113.50**

```
Listening on: [default:]203.0.113.50:8002
Status: listening
Total Bytes Rx: 1230 Tx: 482
Forwarding to: [default:]192.0.0.4:830
Total Bytes Rx: 482 Tx: 1230
Allowed Source Prefixes:
198.51.100.3/32
Connections Max: 1 Active: 1
Source: 198.51.100.3:55002
Listening: connected Bytes Rx: 587 Tx: 621
Forwarding: connected Bytes Rx: 621 Tx: 587
```

user@host> **clear tcp-forwarding statistics listening-port 8002 listening-address 203.0.113.50**

user@host> **show tcp-forwarding status listening-port 8002 listening-address 203.0.113.50**

```
Listening on: [default:]203.0.113.50:8002
Status: listening
Total Bytes Rx: 0 Tx: 0
Forwarding to: [default:]192.0.0.4:830
Total Bytes Rx: 0 Tx: 0
Allowed Source Prefixes:
198.51.100.3/32
```

```

Connections Max: 1 Active: 1
Source: 198.51.100.3:55002
  Listening: connected Bytes Rx: 0 Tx: 0
  Forwarding: connected Bytes Rx: 0 Tx: 0

```

clear tcp-forwarding statistics (Source Address and Port)

The following sample output displays the TCP connection status before and after the connection is cleared.

```
user@host> show tcp-forwarding status
```

```

Listening on: [default:]203.0.113.50:8020
Status: listening
Total Bytes Rx: 292 Tx: 112
Forwarding to: [default:]198.51.100.1:49
Total Bytes Rx: 112 Tx: 292
Allowed Source Prefixes:
192.0.0.1/24
Connections Max: 4 Active: 3
Source: 192.0.0.2:55000
  Listening: connected Bytes Rx: 380 Tx: 223
  Forwarding: connected Bytes Rx: 223 Tx: 380
Source: 192.0.0.3:55000
  Listening: connected Bytes Rx: 855 Tx: 411
  Forwarding: connected Bytes Rx: 411 Tx: 855
Source: 192.0.0.4:56022
  Listening: connected Bytes 642 Tx: 350
  Forwarding: connected Bytes Rx: 350 Tx: 642

```

```
user@host> clear tcp-forwarding statistics source-address 192.0.0.4 source-port 56022
```

```
user@host> show tcp-forwarding status
```

```

Listening on: [default:]203.0.113.50:8020
Status: listening
Total Bytes Rx: 292 Tx: 112
Forwarding to: [default:]198.51.100.1:49
Total Bytes Rx: 112 Tx: 292
Allowed Source Prefixes:
192.0.0.1/24
Connections Max: 4 Active: 3
Source: 192.0.0.2:55000
  Listening: connected Bytes Rx: 380 Tx: 223

```



```
Forwarding: connected Bytes Rx: 223 Tx: 380
Source: 192.0.0.3:55000
Listening: connected Bytes Rx: 855 Tx: 411
Forwarding: connected Bytes Rx: 411 Tx: 855
Source: 192.0.0.4:56022
Listening: connected Bytes 0 Tx: 0
Forwarding: connected Bytes Rx: 0 Tx: 0
```

request interface rebalance (Aggregated Ethernet for Subscriber Management)

Syntax

```
request interface rebalance interface interface-name
```

Release Information

Command introduced in Junos OS Release 11.2.

Description

Manually rebalance the subscribers on an aggregated Ethernet bundle with targeted distribution enabled.

Options

interface-name—Aggregated Ethernet logical interface number.

Required Privilege Level

view

List of Sample Output

[request interface rebalance on page 1284](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
request interface rebalance
```

```
user@host >request interface rebalance interface ae0
```

request network-access aaa subscriber add session-id

Syntax

```
request network-access aaa subscriber add session-id subscriber-session-id service-profile profile-name
```

Release Information

Command introduced in Junos OS Release 11.2.

Description

Locally activate (provision) a dynamic subscriber service for a subscriber who is currently logged in to the network. Starting in Junos OS Release 18.3R1, when the dynamic service profile is configured with the [profile-type remote-device-service](#) statement, the service is provisioned on a remote device by the remote device services manager daemon (rdmd).

Options

profile-name—Name of service-profile to activate.

subscriber-session-id—ID of the subscriber session for which the service will be added.

Required Privilege Level

view

RELATED DOCUMENTATION

CLI-Activated Subscriber Services 20
Local and Remote Service Activation and Deactivation Using the CLI 21
request network-access aaa subscriber delete session-id 1287

List of Sample Output

- [request network-access aaa subscriber add session-id service-profile on page 1286](#)
- [request network-access aaa subscriber add session-id service-profile \(Parameters for Profile on Remote Device\) on page 1286](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request. [Table 54 on page 1285](#) lists possible error messages that might be returned if the service activation fails.

Table 54: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
Command failed: <i>reason</i>	–	–

Table 54: Service Activation/Deactivation Error Messages (*continued*)

Message	Description	Corrective Action
Error: AUTHD ISSU in progress	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.
Provisioning is already active	Remote provisioning by a JSRC server or Gx-plus server is active.	–
Service activation/deactivation already in progress	Another service activation/deactivation operation is currently in progress.	Wait until the active operation completes and then retry the activation/deactivation operation.
Session identifier is not for a subscriber session	The session ID is incorrect.	Verify the correct session ID for the subscriber and then retry the activation/deactivation operation.

Sample Output

```
request network-access aaa subscriber add session-id service-profile
```

```
user@host> request network-access aaa subscriber add session-id 49 service-profile service-bronze
```

```
Successful completion
```

Sample Output

```
request network-access aaa subscriber add session-id service-profile (Parameters for Profile on Remote Device)
```

```
user@host> request network-access aaa subscriber add session-id 131 service-profile
```

```
"upstreamBandwidth(100,100,100)"
```

```
Successful completion
```

request network-access aaa subscriber delete session-id

Syntax

```
request network-access aaa subscriber delete session-id subscriber-session-id service-profile profile-name
```

Release Information

Command introduced in Junos OS Release 11.2.

Description

Deactivate (deprovision) a dynamic subscriber service for a subscriber who is currently logged in to the network. Starting in Junos OS Release 18.3R1, when the dynamic service profile is configured with the [profile-type remote-device-service](#) statement, the service is deprovisioned on a remote device by the remote device services manager daemon (rdmd).

Options

profile-name—Name of the service-profile to deactivate. To deactivate a single instance of a subscriber service that has multiple instances, you can specify the service-profile name and its service parameters.

subscriber-session-id—ID of the subscriber session for which the service will be deleted.

Required Privilege Level

view

RELATED DOCUMENTATION

[CLI-Activated Subscriber Services | 20](#)

[Local and Remote Service Activation and Deactivation Using the CLI | 21](#)

[Deactivating a Single Instance of a Subscriber Service | 28](#)

[Deactivating All Instances of a Subscriber Service | 31](#)

[request network-access aaa subscriber add session-id | 1285](#)

List of Sample Output

[request network-access aaa subscriber delete session-id service-profile on page 1288](#)

[request network-access aaa subscriber delete session-id service-profile \(Deactivating a Single Server Instance\) on page 1288](#)

[request network-access aaa subscriber delete session-id service-profile \(Deactivating All Server Instances\) on page 1289](#)

[request network-access aaa subscriber add session-id service-profile \(Parameters for Profile on Remote Device\) on page 1289](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

[Table 55 on page 1288](#) lists possible error messages that might be returned if the service deactivation fails.

Table 55: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
Command failed: <i>reason</i>	Error condition that caused the command to fail.	Correct the error condition.
Error: AUTHD ISSU in progress	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.
Provisioning is already active	Remote provisioning by a JSRC server or Gx-plus server is active.	Disable provisioning.
Service activation/deactivation already in progress	Another service activation/deactivation operation is currently in progress.	Wait until the active operation completes and then retry the activation/deactivation operation.
Session identifier is not for a subscriber session	The session ID is incorrect.	Verify the correct session ID for the subscriber and then retry the activation/deactivation operation.

Sample Output

request network-access aaa subscriber delete session-id service-profile

```
user@host> request network-access aaa subscriber delete session-id 49 service-profile service-silver
```

```
Successful completion
```

request network-access aaa subscriber delete session-id service-profile (Deactivating a Single Server Instance)

```
user@host> request network-access aaa subscriber delete session-id 6 service-profile
economy-service(up-filter,down-filter)
```

```
Successful completion
```

request network-access aaa subscriber delete session-id service-profile (Deactivating All Server Instances)

user@host> **request network-access aaa subscriber delete session-id 6 service-profile economy-service**

Successful completion

request network-access aaa subscriber add session-id service-profile (Parameters for Profile on Remote Device)

user@host> **request network-access aaa subscriber delete session-id 131 service-profile
"upstreamBandwidth(100,100,100)"**

Successful completion

request network-access aaa subscriber modify session-id

Syntax

```
request network-access aaa subscriber modify session-id subscriber-session-id predefined-variable variable-option
```

Release Information

Command introduced in Junos OS Release 11.2.

Description

Modify a predefined variable that is applied to a subscriber who is currently logged in to the network.

Options

predefined-variable—Name of the predefined variable that you want to modify.

subscriber-session-id—ID of the subscriber session.

variable-option—Name of the variable option that you want to apply to the predefined variable.

Required Privilege Level

view

RELATED DOCUMENTATION

Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers 50
CLI-Activated Subscriber Services 20

List of Sample Output

[request network-access aaa subscriber modify session-id on page 1291](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

[Table 56 on page 1290](#) lists possible messages that might be returned.

Table 56: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
Successful completion	Variable was successfully modified	–
Error: AUTHD ISSU in progress	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.

Sample Output

```
request network-access aaa subscriber modify session-id
```

```
user@host> request network-access aaa subscriber modify session-id 49 junos-cos-traffic-control-profile  
TCP-gold
```

```
Successful completion
```

request network-access aaa subscriber set session-id

Syntax

```
request network-access aaa subscriber set session-id subscriber-session-id provisioning-state none
```

Release Information

Command introduced in Junos OS Release 12.3.

Description

Release control of the PCRF over the specified subscriber session. In response, AAA clears the subscriber's provisioning state and sends a terminated request to the PCRF indicating the subscriber is no longer available.

Options

subscriber-session-id—ID of the subscriber session.

Required Privilege Level

view

RELATED DOCUMENTATION

Disabling PCRF Control of a Subscriber Session
Local and Remote Service Activation and Deactivation Using the CLI 21

List of Sample Output

[request network-access aaa subscriber set session-id on page 1293](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request. [Table 57 on page 1292](#) lists possible error messages that might be returned if the service activation fails.

Table 57: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
Error: AUTHD ISSU in progress	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.

Table 57: Service Activation/Deactivation Error Messages (*continued*)

Message	Description	Corrective Action
Service activation/deactivation already in progress	Another service activation/deactivation operation is currently in progress.	Wait until the active operation completes and then retry the activation/deactivation operation.
Session identifier is not for a subscriber session	The session ID is incorrect.	Verify the correct session ID for the subscriber and then retry the activation/deactivation operation.

Sample Output

```
request network-access aaa subscriber set session-id
```

```
user@host> request network-access aaa subscriber set session-id session-id 49 provisioning-state
none
```

```
Successful completion
```

request services application-identification application

Syntax

```
request services application-identification application <disable | enable> predefined-application-name
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Disable or enable a predefined application signature.

Options

predefined-application-name—Application name; a maximum of up to 31 characters. Predefined applications have the prefix **junos-** to avoid conflict with user-defined ones. Do not name your custom application signature with the **junos** prefix; this prefix is reserved for predefined application signatures.

disable— (Optional) Disable a predefined application signature, initiate signature recompilation, and commit all pending uncompiled signatures to the configuration.

The following conditions apply:

- You cannot disable a predefined application signature that is referenced by an active security policy or custom application signature. First modify or deactivate the policy or custom application signature.
- If you disable an application signature, for example, **junos:HTTP**, that has nested applications, the nested applications are not recognized.

enable—(Optional) Enable a predefined application signature, initiate signature recompilation, and commit all pending uncompiled signatures to the configuration. Include the **no-commit** keyword to defer signature recompilation.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[show services application-identification application](#) | 1469

List of Sample Output

[request services application-identification application disable on page 1295](#)

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request services application-identification application disable

user@host> request services application-identification application disable junos:163

```
Please wait while we are re-compiling signatures ..  
Please wait while we are re-compiling signatures ..  
Please wait while we are re-compiling signatures ..  
Please wait while we are re-compiling signatures ..  
Disable application junos:163 succeed.
```

request services application-identification download

Syntax

```
request services application-identification download <version version-number>;
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Manually download the application package for Junos OS application identification. The application package is extracted from the IDP signature database and contains signature definitions for known applications, such as DNS, Facebook, FTP, Skype, and SNMP.

Options

version version-number—(Optional) Download the specified version of the application package from the Juniper Networks website. If you do not enter a version, the most recent version is downloaded.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[request services application-identification install | 1300](#)

[request services application-identification download status | 1297](#)

List of Sample Output

[request services application-identification download on page 1296](#)

Output Fields

When you enter this command, you are shown the command to use to check the status of your download.

Sample Output

```
request services application-identification download
```

```
user@host> request services application-identification download
```

```
Please use command "request services application-identification download status"
to check status
```

request services application-identification download status

Syntax

```
request services application-identification download status
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Check the download status of the application signature package. The downloaded application package is saved under `/var/db/appid/sec-download/`.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[request services application-identification download](#) | [1296](#)

List of Sample Output

[request services application-identification download status on page 1297](#)

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

```
request services application-identification download status
```

```
user@host> request services application-identifications download status
```

```
Application package 1608 is downloaded successfully.
```

request services application-identification group

Syntax

```
request services application-identification group (copy | disable | enable) predefined-application-group-name
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Copy, disable, or enable a predefined application signature group.

Options

predefined-application-group-name—Identifier for the application group. Maximum length is 32 characters.

copy—Copy the specified predefined application signature group from the database to the configuration and change the name (for example, my:FTP). The ID and order are generated automatically. You can copy the same predefined application signature group only once. You cannot copy duplicate custom signature groups.

NOTE: In configuration mode, if an uncommitted action is pending, the **request services application-identification group copy** command fails.

disable—Disable the specified predefined application signature group.

NOTE: You cannot disable a predefined application signature group that is referenced by an active security policy or custom application signature group. First modify or deactivate the policy or custom application signature group.

enable—Enable the specified predefined application signature group.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

| [show services application-identification group](#) | 1487

List of Sample Output

[request services application-identification group copy on page 1299](#)

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request services application-identification group copy

user@host> request services application-identification group copy junos:SYBASE

```
group 1040 copied successfully.
```

request services application-identification install

Syntax

```
request services application-identification install
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Install the downloaded predefined application signature package.

The install operation fails if any custom application signatures or custom application signature groups have been manually inserted before any predefined application signatures or predefined application signature groups in the Junos OS configuration. Remove any insert-before signatures, then retry the install operation. This command does not display the installation status and only provides an informational message on the types of commands to use to verify the installation status of the application signature package and the protocol bundle.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[request services application-identification download | 1296](#)

[request services application-identification install status | 1302](#)

List of Sample Output

[request services application-identification install on page 1300](#)

Output Fields

When you enter this command, you are shown the command to use to check the status of your installation request.

Sample Output

```
request services application-identification install
```

```
user@host> request services application-identification install
```

Please use command "request services application-identification install status" to check status and use command "request services application-identification proto-bundle-status" to check protocol bundle status

request services application-identification install status

Syntax

```
request services application-identification install status
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display the status of the install operation.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[request services application-identification install](#) | 1300

List of Sample Output

[request services application-identification install status on page 1302](#)

Output Fields

When you enter this command, the system provides feedback on whether your request succeeded or failed.

Sample Output

request services application-identification install status

```
user@host> request services application-identification install status
```

```
Install application package version (1776) succeed.
```

request services application-identification proto-bundle-status

Syntax

```
request services application-identification proto-bundle-status
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display the status of the install operation of the protocol bundle. This command provides feedback on whether your request succeeded or failed.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[request services application-identification install](#) | 1300

List of Sample Output

[request services application-identification proto-bundle-status on page 1303](#)

Output Fields

When you enter this command, the system provides feedback on whether your request succeeded or failed.

Sample Output

```
request services application-identification proto-bundle-status
```

```
user@host> request services application-identification proto-bundle-status
```

```
Protocol Bundle Version (1.30.4-22.005 (build date Jan 17 2014)) and application  
secpack version (2345) is loaded and activated.
```

request services application-identification uninstall

Syntax

```
request services application-identification uninstall
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Uninstall the predefined application package.

The uninstall operation fails if any active security policies, custom application signatures, or custom application signature groups reference predefined application signatures or predefined application signature groups in the Junos OS configuration. This command does not display the uninstallation status and only provides an informational message on the types of commands to use to verify the uninstallation status of the application signature package and the protocol bundle.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[request services application-identification install](#) | [1300](#)

List of Sample Output

[request services application-identification uninstall on page 1304](#)

Output Fields

When you enter this command, you are shown the command to use to check the status of your uninstall request.

Sample Output

request services application-identification uninstall

```
user@host> request services application-identification uninstall
```

```
Please use command "request services application-identification uninstall status"
to check status and use command "request services application-identification
proto-bundle-status" to check protocol bundle status
```

request services application-identification uninstall status

Syntax

```
request services application-identification uninstall status
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display the status of the uninstall operation. This command provides information on whether the uninstall operation succeeded or failed.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

| [request services application-identification uninstall](#) | [1304](#)

List of Sample Output

[request services application-identification uninstall status on page 1305](#)

Output Fields

When you enter this command, the system provides feedback on whether the request succeeded or failed..

Sample Output

```
request services application-identification uninstall status
```

```
user@host> request services application-identification uninstall status
```

```
Uninstall application package version (1776) succeed.
```

request services remote-device-management reconfigure service-device

Syntax

```
request services remote-device-management reconfigure service-device device-name
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Reconfigure a remote device to provision all active subscriber services matching the access domain (list of VLAN ranges and IDs) configured for this remote device with the **vlan-id-list** option at the **[[edit system services remote-device-management service-device *device-name* access-domain]** hierarchy level.

Options

device-name—System-wide name of the remote device that uniquely identifies the device across routing instances.

Required Privilege Level

view

RELATED DOCUMENTATION

[Reconfiguring a Remote Device for RDSM | 652](#)

[request services remote-device-management reload-dictionary | 1308](#)

[Remote Device Services Manager \(RDSM\) Overview | 631](#)

List of Sample Output

[request services remote-device-management reconfigure service-device \(Successful\) on page 1306](#)

[request services remote-device-management reconfigure service-device \(Failed\) on page 1307](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services remote-device-management reconfigure service-device (Successful)

```
user@host> request services remote-device-management reconfigure service-device OLT1
```



```
succeeded
```

request services remote-device-management reconfigure service-device (Failed)

```
user@host> request services remote-device-management reconfigure service-device OLT1
```

```
Service device reconfiguration failed
```

request services remote-device-management reload-dictionary

Syntax

```
request services remote-device-management reload-dictionary absolute file path
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Reload the specified dictionary to the RDSM database. The dictionary is configured with the **dictionary** option at the **[edit system services remote-device-management service-device device-name]** hierarchy level. The reload affects all remote service devices that are configured with this dictionary. When you modify an existing dictionary, this is how you apply the updated file.

NOTE: You cannot modify a dictionary when there is at least one active subscriber service configured on a remote device using that dictionary.

Options

absolute file path—Absolute file path for the vendor-specific dictionary that defines the set of NETCONF XML protocol commands required to provision, deprovision, and roll back a subscriber service on the remote device. The dictionary is stored on the BNG. An example absolute path is ***/var/home/dict/remote-device.xml***.

Required Privilege Level

view

RELATED DOCUMENTATION

[Reloading a Dictionary File for RDSM | 653](#)

[request services remote-device-management reconfigure service-device | 1306](#)

[Remote Device Services Manager \(RDSM\) Overview | 631](#)

List of Sample Output

[request services remote-device-management reload-dictionary \(Successful\) on page 1309](#)

[request services remote-device-management reload-dictionary \(Failed\) on page 1309](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services remote-device-management reload-dictionary (Successful)

```
user@host> request services remote-device-management reload-dictionary  
/var/home/dict/vendor-1-dictionary.xml
```

```
Dictionary reloaded successfully
```

request services remote-device-management reload-dictionary (Failed)

```
user@host> request services remote-device-management reload-dictionary  
/var/home/dict/vendor-2-dictionary.xml
```

```
Dictionary reload failed
```

show class-of-service

Syntax

```
show class-of-service
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entire class-of-service (CoS) configuration, including system-chosen defaults. Executing this command is equivalent to executing all **show class-of-service** commands in succession.

Options

This command has no options.

Required Privilege Level

view

List of Sample Output

[show class-of-service on page 1310](#)

Output Fields

See the output field descriptions for the commands.

Sample Output

show class-of-service

```
user@host> show class-of-service
```

```
Forwarding class           Queue
best-effort                0
expedited-forwarding       1
assured-forwarding         2
network-control            3
Code point type: dscp
Alias      Bit pattern
af11       001010
af12       001100
af13       001110
```

```

...
Code point type: dscp-ipv6
  Alias          Bit pattern
  af11           001010
  af12           001100
  af13           001110
...
Code point type: exp
  Alias          Bit pattern
  af11           100
  af12           101
  be             000
...
Code point type: ieee-802.1
  Alias          Bit pattern
  af11           100
  af12           101
  be             000
...
Classifier: dscp-default, Code point type: dscp, Index: 6
  Code point      Forwarding class      Loss priority
  000000          best-effort            low
  000001          best-effort            low
  000010          best-effort            low
....
Classifier: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 7
  Code point      Forwarding class      Loss priority
  000000          best-effort            low
  000001          best-effort            low
  000010          best-effort            low
...
Loss-priority-map: frame-relay-de-default, Code point type: frame-relay-de, Index:
12
  Code point      Loss priority
  0               low
  1               high

Rewrite rule: dscp-default, Code point type: dscp, Index: 23
  Forwarding class      Loss priority      Code point
  best-effort           low                000000
  best-effort           high               000000
  expedited-forwarding  low                101110
...
Rewrite rule: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 24

```

```

Forwarding class          Loss priority    Code point
best-effort               low          000000
best-effort               high         000000
...
....
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
      100          100

Scheduler map: <default>, Index: 2

Scheduler: <default-be>, Forwarding class: best-effort, Index: 16
  Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent, Priority:
low
  Drop profiles:
    Loss priority  Protocol    Index    Name
    Low           any         1        <default-drop-profile>
    Medium low    any         1        <default-drop-profile>
    Medium high   any         1        <default-drop-profile>
    High          any         1        <default-drop-profile>
...
Physical interface: fe-0/0/0, Index: 137
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2

Logical interface: fe-0/0/0.0, Index: 69
  Object          Name                      Type          Index
  Adaptive-shaper  fr-shaper                               35320
  Classifier       ipprec-compatibility  ip            11

Physical interface: fe-0/0/1, Index: 138
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2
...

```

show class-of-service adjustment-control-profile

Syntax

```
show class-of-service adjustment-control-profile
<profile-name>
```

Release Information

Command introduced in Junos OS Release 13.1 for MX Series Routers.

Description

For MPC/MIC interfaces only, display the adjustment control profiles.

Options

none—Display all profiles.

profile-name—(Optional) Display information about a single profile.

Required Privilege Level

view

RELATED DOCUMENTATION

Verifying the CoS Adjustment Control Profile Configuration | 188

List of Sample Output

[show class-of-service adjustment-control-profile on page 1314](#)

Output Fields

[Table 58 on page 1313](#) describes the output fields for the **show class-of-service adjustment-control-profile** command. Output fields are listed in the approximate order in which they appear.

Table 58: show class-of-service adjustment-control-profile Output Fields

Field Name	Field Description
Name	<p>Name of the adjusting application. Possible values:</p> <ul style="list-style-type: none"> • RADIUS-CoA—RADIUS CoA application. • ANCP—ANCP application. • PPPoE IA tags—PPPoE IA tag application. • DHCP tags—DHCP application.

Table 58: show class-of-service adjustment-control-profile Output Fields (*continued*)

Field Name	Field Description
Priority	<p>Priority of the adjusting application. Possible values are 1 through 10; 1 is the highest priority.</p> <p>The lower the priority value, the higher the priority</p>
Algorithm	<p>Algorithm the adjusting application uses to make adjustments.</p> <ul style="list-style-type: none"> • adjust-never—Never 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.

Sample Output

show class-of-service adjustment-control-profile

```
user@host> show class-of-service adjustment-control-profile
```

```

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
  name: DHCP tags, priority: 2, algorithm: less

```


show class-of-service interface

Syntax

```
show class-of-service interface  
<comprehensive | detail> <interface-name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Forwarding class map information added in Junos OS Release 9.4.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport routers.

Command introduced in Junos OS Release 12.2 for the ACX Series Universal Metro routers.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Options **detail** and **comprehensive** introduced in Junos OS Release 11.4.

Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description

Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.

NOTE: On routing platforms with dual Routing Engines, running this command on the backup Routing Engine, with or without any of the available options, is not supported and produces the following error message:

error: the class-of-service subsystem is not running

Options

none—Display CoS associations for all physical and logical interfaces.

comprehensive—(M Series, MX Series, and T Series routers) (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.

detail—(M Series, MX Series, and T Series routers) (Optional) Display QoS and CoS information based on the interface.

If the **interface** *interface-name* is a physical interface, the output includes:

- Brief QoS information about the physical interface
- Brief QoS information about the logical interface
- CoS information about the physical interface
- Brief information about filters or policers of the logical interface
- Brief CoS information about the logical interface

If the **interface** *interface-name* is a logical interface, the output includes:

- Brief QoS information about the logical interface
- Information about filters or policers for the logical interface
- CoS information about the logical interface

interface-name—(Optional) Display class-of-service (CoS) associations for the specified interface.

none—Display CoS associations for all physical and logical interfaces.

NOTE: ACX5000 routers do not support classification on logical interfaces and therefore do not show CoS associations for logical interfaces with this command.

Required Privilege Level

view

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

List of Sample Output

[show class-of-service interface \(Physical\) on page 1332](#)

[show class-of-service interface \(Logical\) on page 1332](#)

[show class-of-service interface \(Gigabit Ethernet\) on page 1332](#)

[show class-of-service interface \(ANCP\) on page 1333](#)

[show class-of-service interface \(PPPoE Interface\) on page 1333](#)

[show class-of-service interface \(DHCP Interface\) on page 1333](#)

[show class-of-service interface \(T4000 Routers with Type 5 FPCs\) on page 1334](#)

[show class-of-service interface detail on page 1334](#)

[show class-of-service interface comprehensive on page 1335](#)

[show class-of-service interface \(ACX Series Routers\) on page 1350](#)

[show class-of-service interface \(PPPoE Subscriber Interface for Enhanced Subscriber Management\) on page 1353](#)

Output Fields

[Table 59 on page 1317](#) describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

Table 59: show class-of-service interface Output Fields

Field Name	Field Description
Physical interface	Name of a physical interface.
Index	Index of this interface or the internal index of this object. (Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles and dynamic scheduler maps are larger for enhanced subscriber management than they are for legacy subscriber management.
Dedicated Queues	Status of dedicated queues configured on an interface. Supported only on Trio MPC/MIC interfaces on MX Series routers. (Enhanced subscriber management for MX-Series routers) This field is not displayed for enhanced subscriber management.
Maximum usable queues	Number of queues you can configure on the interface.
Maximum usable queues	Maximum number of queues you can use.
Total non-default queues created	Number of queues created in addition to the default queues. Supported only on Trio MPC/MIC interfaces on MX Series routers. (Enhanced subscriber management for MX Series routers) This field is not displayed for enhanced subscriber management.
Rewrite Input IEEE Code-point	(QFX3500 switches only) IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.
Shaping rate	Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the Shaping rate field is displayed for either the physical interface or the logical interface.

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Scheduler map	Name of the output scheduler map associated with this interface. (Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, SMAP-1_UID1002) instead of with a subscriber interface.
Scheduler map forwarding class sets	(QFX Series only) Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.
Input shaping rate	For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.
Input scheduler map	For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.
Chassis scheduler map	Name of the scheduler map associated with the packet forwarding component queues.
Rewrite	Name and type of the rewrite rules associated with this interface.
Traffic-control-profile	Name of the associated traffic control profile. (Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, TC_PROF_100_199_SERIES_UID1006) instead of with a subscriber interface.
Classifier	Name and type of classifiers associated with this interface.
Forwarding-class-map	Name of the forwarding map associated with this interface.
Congestion-notification	(QFX Series and EX4600 switches only) Congestion notification state, enabled or disabled .
Logical interface	Name of a logical interface.
Object	Category of an object: Classifier , Fragmentation-map (for LSQ interfaces only), Scheduler-map , Rewrite , Translation Table (for IQE PICs only), or traffic-class-map (for T4000 routers with Type 5 FPCs).
Name	Name of an object.
Type	Type of an object: dscp , dscp-ipv6 , exp , ieee-802.1 , ip , inet-precedence , or ieee-802.1ad (for traffic class map on T4000 routers with Type 5 FPCs)..
Link-level type	Encapsulation on the physical interface.

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback.
Source filtering	Whether source filtering is enabled or disabled.
Flow control	Whether flow control is enabled or disabled.
Auto-negotiation	(Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status. <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline.
Device flags	The Device flags field provides information about the physical device and displays one or more of the following values: <ul style="list-style-type: none"> • Down—Device has been administratively disabled. • Hear-Own-Xmit—Device receives its own transmissions. • Link-Layer-Down—The link-layer protocol has failed to connect with the remote endpoint. • Loopback—Device is in physical loopback. • Loop-Detected—The link layer has received frames that it sent, thereby detecting a physical loopback. • No-Carrier—On media that support carrier recognition, no carrier is currently detected. • No-Multicast—Device does not support multicast traffic. • Present—Device is physically present and recognized. • Promiscuous—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. • Quench—Transmission on the device is quenched because the output buffer is overflowing. • Recv-All-Multicasts—Device is in multicast promiscuous mode and therefore provides no multicast filtering. • Running—Device is active and enabled.

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Interface flags	<p>The Interface flags field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> • Admin-Test—Interface is in test mode and some sanity checking, such as loop detection, is disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Hardware-Down—Interface is nonfunctional or incorrectly connected. • Link-Layer-Down—Interface keepalives have indicated that the link is incomplete. • No-Multicast—Interface does not support multicast traffic. • No-receive No-transmit—Passive monitor mode is configured on the interface. • Point-To-Point—Interface is point-to-point. • Pop all MPLS labels from packets of depth—MPLS labels are removed as packets arrive on an interface that has the pop-all-labels statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> • 1—Takes effect for incoming packets with one label only. • 2—Takes effect for incoming packets with two labels only. • [1 2]—Takes effect for incoming packets with either one or two labels. • Promiscuous—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses. • Recv-All-Multicasts—Interface is in multicast promiscuous mode and provides no multicast filtering. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Flags	<p>The Logical interface flags field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> • ACFC Encapsulation—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer). • Device-down—Device has been administratively disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Clear-DF-Bit—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit. • Hardware-Down—Interface protocol initialization failed to complete successfully. • PFC—Protocol field compression is enabled for the PPP session. • Point-To-Point—Interface is point-to-point. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.
Encapsulation	Encapsulation on the logical interface.
Admin	Administrative state of the interface (Up or Down)
Link	Status of physical link (Up or Down).
Proto	Protocol configured on the interface.
Input Filter	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
Output Filter	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Link flags	<p>Provides information about the physical link and displays one or more of the following values:</p> <ul style="list-style-type: none"> • ACFC—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option. • Give-Up—Link protocol does not continue connection attempts after repeated failures. • Loose-LCP—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational. • Loose-LMI—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational. • Loose-NCP—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational. • Keepalives—Link protocol keepalives are enabled. • No-Keepalives—Link protocol keepalives are disabled. • PFC—Protocol field compression is configured. The PPP session negotiates the PFC option.
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
CoS queues	Number of CoS queues configured.
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .
Statistics last cleared	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface.
Exclude Overhead Bytes	<p>Exclude the counting of overhead bytes from aggregate queue statistics.</p> <ul style="list-style-type: none"> • Disabled—Default configuration. Includes the counting of overhead bytes in aggregate queue statistics. • Enabled—Excludes the counting of overhead bytes from aggregate queue statistics for just the physical interface. • Enabled for hierarchy—Excludes the counting of overhead bytes from aggregate queue statistics for the physical interface as well as all child interfaces, including logical interfaces and interface sets.
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Input errors	<p>Input errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Bucket Drops—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the <code>ignore-l3-incompletes</code> statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • HS link FIFO overflows—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Output errors	<p>Output errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Drops field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • HS link FIFO underflows—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeds the MTU of the interface.
Egress queues	Total number of egress Maximum usable queues on the specified interface.
Queue counters	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Dropped packets field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
SONET alarms SONET defects	<p>(SONET) SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: SONET PHY, SONET section, SONET line, and SONET path.</p>

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
SONET PHY	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET PHY field has the following subfields:</p> <ul style="list-style-type: none"> • PLL Lock—Phase-locked loop • PHY Light—Loss of optical signal
SONET section	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET section field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOS—Loss of signal • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section)

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
SONET line	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET line field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line)

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
SONET path	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET path field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • ES-PFE—Errored seconds (far-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path)
Received SONET overhead Transmitted SONET overhead	<p>Values of the received and transmitted SONET overhead:</p> <ul style="list-style-type: none"> • C2—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • F1—Section user channel byte. This byte is set aside for the purposes of users. • K1 and K2—These bytes are allocated for APS signaling for the protection of the multiplex section. • J0—Section trace. This byte is defined for STS-1 number 1 of an STS-<i>N</i> signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter. • S1—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-<i>N</i> signal. • Z3 and Z4—Allocated for future use.

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Received path trace Transmitted path trace	SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.
HDLC configuration	Information about the HDLC configuration. <ul style="list-style-type: none"> • Policing bucket—Configured state of the receiving policer. • Shaping bucket—Configured state of the transmitting shaper. • Giant threshold—Giant threshold programmed into the hardware. • Runt threshold—Runt threshold programmed into the hardware.
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> • Destination slot—FPC slot number. • PLP byte—Packet Level Protocol byte.
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.
Forwarding classes	Total number of forwarding classes supported on the specified interface.
Egress queues	Total number of egress Maximum usable queues on the specified interface.
Queue	Queue number.
Forwarding classes	Forwarding class name.

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Queued Packets	Number of packets queued to this queue.
Queued Bytes	Number of bytes queued to this queue. The byte counts vary by PIC type.
Transmitted Packets	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the Packet Forwarding Engine Chassis Queues field) shows the prefragmentation values.
Transmitted Bytes	Number of bytes transmitted by this queue. The byte counts vary by PIC type.
Tail-dropped packets	Number of packets dropped because of tail drop.
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> • (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> • Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. • Low, TCP—Number of low-loss priority TCP packets dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. • High, TCP—Number of high-loss priority TCP packets dropped because of RED. • (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> • Low—Number of low-loss priority packets dropped because of RED. • Medium-low—Number of medium-low loss priority packets dropped because of RED. • Medium-high—Number of medium-high loss priority packets dropped because of RED. • High—Number of high-loss priority packets dropped because of RED. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by PIC type.</p> <ul style="list-style-type: none"> • (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> • Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. • Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. • High, TCP—Number of high-loss priority TCP bytes dropped because of RED. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
Transmit rate	Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.
Rate Limit	<p>Rate limiting configuration of the queue. Possible values are :</p> <ul style="list-style-type: none"> • None—No rate limit. • exact—Queue transmits at the configured rate.
Buffer size	Delay buffer size in the queue.
Priority	Scheduling priority configured as low or high .
Excess Priority	Priority of the excess bandwidth traffic on a scheduler: low , medium-low , medium-high , high , or none .
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> • Loss priority—Packet loss priority for drop profile assignment. • Protocol—Transport protocol for drop profile assignment. • Index—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • Name—Name of the drop profile. • Type—Type of the drop profile: discrete or interpolated. • Fill Level—Percentage fullness of a queue. • Drop probability—Drop probability at this fill level.
Excess Priority	Priority of the excess bandwidth traffic on a scheduler.

Table 59: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> • Loss priority—Packet loss priority for drop profile assignment. • Protocol—Transport protocol for drop profile assignment. • Index—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • Name—Name of the drop profile. • Type—Type of the drop profile: discrete or interpolated. • Fill Level—Percentage fullness of a queue. • Drop probability—Drop probability at this fill level.
Adjustment information	<p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> • Adjusting application—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> • The adjusting application can appear as ancp LS-0, which is the Junos OS Access Node Control Profile process (ancpd) that performs shaping-rate adjustments on schedule nodes. • The adjusting application can appear as DHCP, which adjusts the shaping-rate and overhead-accounting class-of-service attributes based on DSL Forum VSA conveyed in DHCP option 82, suboption 9 (Vendor Specific Information). The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode). • The adjusting application can also appear as pppoe, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode). • Adjustment type—Type of adjustment: absolute or delta. • Configured shaping rate—Shaping rate configured for the scheduler node or queue. • Adjustment value—Value of adjusted shaping rate. • Adjustment target—Level of shaping-rate adjustment performed: node or queue. • Adjustment overhead-accounting mode—Configured shaping mode: frame or cell. • Adjustment overhead bytes—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS. • Adjustment target—Level of shaping-rate adjustment performed: node or queue. • Adjustment multicast index—

Sample Output

show class-of-service interface (Physical)

user@host> show class-of-service interface so-0/2/3

```
Physical interface: so-0/2/3, Index: 135
Maximum usable queues: 8, Queues in use: 4
Total non-default queues created: 4
  Scheduler map: <default>, Index: 2032638653

Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
  Shaping rate: 32000
  Object          Name                Type          Index
  Scheduler-map   <default>           27
  Rewrite         exp-default         exp           21
  Classifier      exp-default         exp           5
  Classifier      ipprec-compatibility ip             8
  Forwarding-class-map exp-default         exp           5
```

show class-of-service interface (Logical)

user@host> show class-of-service interface so-0/2/3.0

```
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
  Shaping rate: 32000
  Object          Name                Type          Index
  Scheduler-map   <default>           27
  Rewrite         exp-default         exp           21
  Classifier      exp-default         exp           5
  Classifier      ipprec-compatibility ip             8
  Forwarding-class-map exp-default         exp           5
```

show class-of-service interface (Gigabit Ethernet)

user@host> show class-of-service interface ge-6/2/0

```
Physical interface: ge-6/2/0, Index: 175
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Input scheduler map: <default>, Index: 3
  Chassis scheduler map: <default-chassis>, Index: 4
```

show class-of-service interface (ANCP)

```
user@host> show class-of-service interface pp0.1073741842
```

```
Logical interface: pp0.1073741842, Index: 341
```

Object	Name	Type	Index
Traffic-control-profile	TCP-CVLAN	Output	12408
Classifier	dscp-ipv6-compatibility	dscp-ipv6	9
Classifier	ipprec-compatibility	ip	13

```

Adjusting application: ancp LS-0
Adjustment type: absolute
Configured shaping rate: 4000000
Adjustment value: 11228000
Adjustment overhead-accounting mode: Frame Mode
Adjustment overhead bytes: 50
Adjustment target: node

```

show class-of-service interface (PPPoE Interface)

```
user@host> show class-of-service interface pp0.1
```

```
Logical interface: pp0.1, Index: 85
```

Object	Name	Type	Index
Traffic-control-profile	tcp-pppoe.o.pp0.1	Output	2726446535
Classifier	ipprec-compatibility	ip	13

```

Adjusting application: PPPoE
Adjustment type: absolute
Adjustment value: 5000000
Adjustment overhead-accounting mode: cell
Adjustment target: node

```

show class-of-service interface (DHCP Interface)

```
user@host> show class-of-service interface demux0.1
```

```
Logical interface: pp0.1, Index: 85
```

Object	Name	Type	Index
Traffic-control-profile	tcp-dhcp.o.demux0.1	Output	2726446535
Classifier	ipprec-compatibility	ip	13

```

Adjusting application: DHCP
Adjustment type: absolute

```

```
Adjustment value: 5000000
Adjustment overhead-accounting mode: cell
Adjustment target: node
```

show class-of-service interface (T4000 Routers with Type 5 FPCs)

user@host> **show class-of-service interface xe-4/0/0**

```
Physical interface: xe-4/0/0, Index: 153
  Maximum usable queues: 8, Queues in use: 4
  Shaping rate: 5000000000 bps
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: xe-4/0/0.0, Index: 77
    Object          Name          Type
Index
  Classifier        ipprec-compatibility ip
13
```

show class-of-service interface detail

user@host> **show class-of-service interface ge-0/3/0 detail**

```
Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, Loopback: Disabled, Source
  filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote
  fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

Physical interface: ge-0/3/0, Index: 138
  Maximum usable queues: 4, Queues in use: 5
  Shaping rate: 50000 bps
  Scheduler map: interface-scheduler-map, Index: 58414
  Input shaping rate: 10000 bps
  Input scheduler map: scheduler-map, Index: 15103
  Chassis scheduler map: <default-chassis>, Index: 4
  Congestion-notification: Disabled

Logical interface ge-0/3/0.0
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
```

```

    inet
    mpls
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/3/0.0     up    up    inet
               mpls

Interface      Admin Link Proto Input Policer      Output Policer
ge-0/3/0.0     up    up    inet
               mpls

Logical interface: ge-0/3/0.0, Index: 68
Object          Name                      Type                      Index
Rewrite         exp-default              exp (mpls-any)           33
Classifier       exp-default              exp                       10
Classifier       ipprec-compatibility     ip                        13

Logical interface ge-0/3/0.1
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2
  inet
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/3/0.1     up    up    inet
Interface      Admin Link Proto Input Policer      Output Policer
ge-0/3/0.1     up    up    inet

Logical interface: ge-0/3/0.1, Index: 69
Object          Name                      Type                      Index
Classifier       ipprec-compatibility     ip                        13

```

show class-of-service interface comprehensive

user@host> **show class-of-service interface ge-0/3/0 comprehensive**

```

Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 601, Generation: 141
  Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
  control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Schedulers     : 256
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d

```

```

Last flapped   : 2010-09-07 06:35:22 PDT (15:14:42 ago)
Statistics last cleared: Never   Exclude Overhead Bytes: Disabled

Traffic statistics:
  Input bytes   :                0                0 bps
  Output bytes  :                0                0 bps
  Input packets:                0                0 pps
  Output packets:               0                0 pps
IPv6 total statistics:
  Input bytes   :                0
  Output bytes  :                0
  Input packets:                0
  Output packets:               0
Ingress traffic statistics at Packet Forwarding Engine:
  Input bytes   :                0                0 bps
  Input packets:                0                0 pps
  Drop bytes    :                0                0 bps
  Drop packets  :                0                0 pps
Label-switched interface (LSI) traffic statistics:
  Input bytes   :                0                0 bps
  Input packets:                0                0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
  Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 5 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 af3              0                0                0
  1 af2              0                0                0
  2 ef2              0                0                0
  3 ef1              0                0                0
Egress queues: 4 supported, 5 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 af3              0                0                0
  1 af2              0                0                0
  2 ef2              0                0                0
  3 ef1              0                0                0
Active alarms   : None
Active defects  : None
MAC statistics:
  Total octets      Receive          Transmit
  Total packets      0                0

```

```

Unicast packets                0                0
Broadcast packets              0                0
Multicast packets              0                0
CRC/Align errors               0                0
FIFO errors                    0                0
MAC control frames             0                0
MAC pause frames               0                0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0
Filter statistics:
  Input packet count            0
  Input packet rejects          0
  Input DA rejects              0
  Input SA rejects              0
  Output packet count           0
  Output packet pad count       0
  Output packet error count     0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
OK
  Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue           Bandwidth           Buffer Priority
Limit
                                %           bps           %           usec
    2 ef2                      39          19500         0           120      high
none
  Direction : Input
  CoS transmit queue           Bandwidth           Buffer Priority
Limit
                                %           bps           %           usec
    0 af3                      30           3000        45            0      low
none

```

```

Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 601
Forwarding classes: 16 supported, 5 in use
Ingress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps

```


Forwarding classes: 16 supported, 5 in use

Egress queues: 4 supported, 5 in use

Queue: 0, Forwarding classes: af3

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : Not Available

RL-dropped packets	:	0	0 pps
--------------------	---	---	-------

RL-dropped bytes	:	0	0 bps
------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Queue: 1, Forwarding classes: af2

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : Not Available

RL-dropped packets	:	0	0 pps
--------------------	---	---	-------

RL-dropped bytes	:	0	0 bps
------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Queue: 2, Forwarding classes: ef2

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : Not Available

RL-dropped packets	:	0	0 pps
--------------------	---	---	-------

RL-dropped bytes	:	0	0 bps
------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Queue: 3, Forwarding classes: ef1

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
Tail-dropped packets	:	Not Available	
RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 5 in use

Queue: 0, Forwarding classes: af3

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	Not Available	
RED-dropped bytes	:	Not Available	

Queue: 1, Forwarding classes: af2

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	Not Available	
RED-dropped bytes	:	Not Available	

Queue: 2, Forwarding classes: ef2

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	Not Available	
RED-dropped bytes	:	Not Available	

Queue: 3, Forwarding classes: ef1

Queued:

Packets	:	108546	0 pps
Bytes	:	12754752	376 bps

Transmitted:

```

Packets          :          108546          0 pps
Bytes            :          12754752        376 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets : Not Available
RED-dropped bytes  : Not Available

```

```

Physical interface: ge-0/3/0, Index: 138
Maximum usable queues: 4, Queues in use: 5
Shaping rate: 50000 bps

```

```
Scheduler map: interface-scheduler-map, Index: 58414
```

```

Scheduler: ef2, Forwarding class: ef2, Index: 39155
  Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer Limit:
  none, Priority: high
  Excess Priority: unspecified
  Drop profiles:
    Loss priority  Protocol      Index      Name
    Low           any           1          < default-drop-profile>
    Medium low    any           1          < default-drop-profile>
    Medium high   any           1          < default-drop-profile>
    High          any           1          < default-drop-profile>
  Drop profile: < default-drop-profile>, Type: discrete, Index: 1
    Fill level    Drop probability
    100           100
  Drop profile: < default-drop-profile>, Type: discrete, Index: 1
    Fill level    Drop probability
    100           100
  Drop profile: < default-drop-profile>, Type: discrete, Index: 1
    Fill level    Drop probability
    100           100
  Drop profile: < default-drop-profile>, Type: discrete, Index: 1
    Fill level    Drop probability
    100           100
  Input shaping rate: 10000 bps
  Input scheduler map: scheduler-map

```

```
Scheduler map: scheduler-map, Index: 15103
```

```

Scheduler: af3, Forwarding class: af3, Index: 35058
  Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer
  Limit: none, Priority: low
  Excess Priority: unspecified
  Drop profiles:

```

Loss priority	Protocol	Index	Name
Low	any	40582	green
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	18928	yellow

Drop profile: green, Type: discrete, Index: 40582

Fill level	Drop probability
50	0
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: yellow, Type: discrete, Index: 18928

Fill level	Drop probability
50	0
100	100

Chassis scheduler map: < default-drop-profile>

Scheduler map: < default-drop-profile>, Index: 4

Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25

Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none, Priority: low

Excess Priority: low

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	< default-drop-profile>
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	1	< default-drop-profile>

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

```

Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25
  Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
  Excess Priority: low
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           any       1      < default-drop-profile>
    Medium low    any       1      < default-drop-profile>
    Medium high   any       1      < default-drop-profile>
    High          any       1      < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100

Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25
  Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
  Excess Priority: low
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           any       1      < default-drop-profile>
    Medium low    any       1      < default-drop-profile>
    Medium high   any       1      < default-drop-profile>
    High          any       1      < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level  Drop probability
    100        100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1

```

```

Fill level      Drop probability
      100              100

Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
  Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
  Excess Priority: low
  Drop profiles:
    Loss priority  Protocol      Index      Name
    Low           any           1          < default-drop-profile>
    Medium low    any           1          < default-drop-profile>
    Medium high   any           1          < default-drop-profile>
    High          any           1          < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
  Fill level      Drop probability
      100              100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level      Drop probability
      100              100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level      Drop probability
      100              100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level      Drop probability
      100              100
  Congestion-notification: Disabled

Forwarding class          ID      Queue  Restricted queue  Fabric
priority Policing priority
af3                      0      0      0              low
      normal
af2                      1      1      1              low
      normal
ef2                      2      2      2              high
      normal
ef1                      3      3      3              high
      normal
af1                      4      4      0              low
      normal

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
  Traffic statistics:
    Input bytes :      0
    Output bytes :     0

```

```

    Input  packets:                0
    Output packets:                0
Local statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets:                0
    Output packets:                0
Transit statistics:
    Input  bytes   :                0          0 bps
    Output bytes   :                0          0 bps
    Input  packets:                0          0 pps
    Output packets:                0          0 pps
Protocol inet, MTU: 1500, Generation: 172, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Input Filters: filter-in-ge-0/3/0.0-i,
    Policer: Input: pl-ge-0/3/0.0-inet-i
Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0
    Flags: Is-Primary
    Output Filters: exp-filter,,,,,

```

Logical interface ge-1/2/0.0 (Index 347) (SNMP ifIndex 638) (Generation 156)

Forwarding class	ID	Queue	Restricted queue	Fabric priority	Policing priority
SPU priority					
best-effort	0	0	0	low	normal
low					

Aggregate Forwarding-class statistics per forwarding-class

Aggregate Forwarding-class statistics:

Forwarding-class statistics:

Forwarding-class best-effort statistics:

```

    Input unicast bytes:    0
    Output unicast bytes:   0
    Input unicast packets:  0
    Output unicast packets: 0

```

```

    Input multicast bytes:   0
    Output multicast bytes:  0
    Input multicast packets: 0
    Output multicast packets: 0

```

Forwarding-class expedited-forwarding statistics:

```

    Input unicast bytes:    0

```

```
Output unicast bytes:      0
Input unicast packets:    0
Output unicast packets:   0
```

```
Input multicast bytes:     0
Output multicast bytes:    0
Input multicast packets:   0
Output multicast packets:  0
```

IPv4 protocol forwarding-class statistics:

Forwarding-class statistics:

Forwarding-class best-effort statistics:

```
Input unicast bytes:      0
Output unicast bytes:     0
Input unicast packets:    0
Output unicast packets:   0
```

```
Input multicast bytes:    0
Output multicast bytes:   0
Input multicast packets:  0
Output multicast packets: 0
```

Forwarding-class expedited-forwarding statistics:

```
Input unicast bytes:      0
Output unicast bytes:     0
Input unicast packets:    0
Output unicast packets:   0
```

```
Input multicast bytes:    0
Output multicast bytes:   0
Input multicast packets:  0
Output multicast packets: 0
```

IPv6 protocol forwarding-class statistics:

Forwarding-class statistics:

Forwarding-class best-effort statistics:

```
Input unicast bytes:      0
Output unicast bytes:     0
Input unicast packets:    0
Output unicast packets:   0
```

```
Input multicast bytes:    0
```



```

Output multicast bytes:    0
Input multicast packets:  0
Output multicast packets: 0

```

Forwarding-class expedited-forwarding statistics:

```

Input unicast bytes:      0
Output unicast bytes:     0
Input unicast packets:    0
Output unicast packets:   0

```

```

Input multicast bytes:    0
Output multicast bytes:   0
Input multicast packets:  0
Output multicast packets: 0

```

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)

```

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
Input packets : 0
Output packets: 0

```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/3/0.0	up	up	inet	filter-in-ge-0/3/0.0-i	
			mpls		exp-filter

Interface	Admin	Link	Proto	Input Policer	Output Policer
ge-0/3/0.0	up	up			
			inet	p1-ge-0/3/0.0-inet-i	
			mpls		

Filter: filter-in-ge-0/3/0.0-i

Counters:

Name	Bytes	Packets
count-filter-in-ge-0/3/0.0-i	0	0

Filter: exp-filter

Counters:

Name	Bytes	Packets
count-exp-seven-match	0	0
count-exp-zero-match	0	0

Policers:

Name	Packets
p1-ge-0/3/0.0-inet-i	0

Logical interface: ge-0/3/0.0, Index: 68

Object	Name	Type	Index
Rewrite	exp-default	exp (mpls-any)	33

Rewrite rule: exp-default, Code point type: exp, Index: 33

Forwarding class	Loss priority	Code point
af3	low	000
af3	high	001
af2	low	010
af2	high	011
ef2	low	100
ef2	high	101
ef1	low	110
ef1	high	111

Object	Name	Type	Index
Classifier	exp-default	exp	10

Classifier: exp-default, Code point type: exp, Index: 10

Code point	Forwarding class	Loss priority
000	af3	low
001	af3	high
010	af2	low
011	af2	high
100	ef2	low
101	ef2	high
110	ef1	low
111	ef1	high

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

Code point	Forwarding class	Loss priority
000	af3	low
001	af3	high
010	af3	low
011	af3	high
100	af3	low
101	af3	high
110	ef1	low
111	ef1	high

Forwarding class	ID	Queue	Restricted queue	Fabric
priority				
Policing priority				
af3	0	0	0	low
normal				
af2	1	1	1	low

	normal				
ef2		2	2	2	high
	normal				
ef1		3	3	3	high
	normal				
af1		4	4	0	low
	normal				

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)

Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2

Traffic statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Local statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Transit statistics:

Input bytes : 0 0 bps

Output bytes : 0 0 bps

Input packets: 0 0 pps

Output packets: 0 0 pps

Protocol inet, MTU: 1500, Generation: 174, Route table: 0

Flags: Sendbroadcast-pkt-to-re

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)

Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2

Input packets : 0

Output packets: 0

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/3/0.1	up	up	mpls		
Interface	Admin	Link	Proto	Input Policer	Output Policer
ge-0/3/0.1	up	up			
			mpls		

Logical interface: ge-0/3/0.1, Index: 69

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

Code point	Forwarding class	Loss priority
000	af3	low
001	af3	high
010	af3	low
011	af3	high
100	af3	low
101	af3	high
110	ef1	low
111	ef1	high

Forwarding class	ID	Queue	Restricted queue	Fabric
priority Policing priority				
af3	0	0	0	low
normal				
af2	1	1	1	low
normal				
ef2	2	2	2	high
normal				
ef1	3	3	3	high
normal				
af1	4	4	0	low
normal				

show class-of-service interface (ACX Series Routers)

user@host-g11# show class-of-service interface

```
Physical interface: at-0/0/0, Index: 130
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

Logical interface: at-0/0/0.0, Index: 69

Logical interface: at-0/0/0.32767, Index: 70

Physical interface: at-0/0/1, Index: 133
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

Logical interface: at-0/0/1.0, Index: 71
```

Logical interface: at-0/0/1.32767, Index: 72

Physical interface: ge-0/1/0, Index: 146

Maximum usable queues: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Rewrite	dscp-default	dscp	31
Classifier	d1	dscp	11331
Classifier	ci	ieee8021p	583

Logical interface: ge-0/1/0.0, Index: 73

Object	Name	Type	Index
Rewrite	custom-exp	exp (mpls-any)	46413

Logical interface: ge-0/1/0.1, Index: 74

Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147

Maximum usable queues: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: ge-0/1/1.0, Index: 76

Physical interface: ge-0/1/2, Index: 148

Maximum usable queues: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Rewrite	ri	ieee8021p (outer)	35392
Classifier	ci	ieee8021p	583

Physical interface: ge-0/1/3, Index: 149

Maximum usable queues: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

```

    Logical interface: ge-0/1/3.0, Index: 77
Object      Name      Type      Index
Rewrite     custom-exp2    exp (mpls-any)  53581

Physical interface: ge-0/1/4, Index: 150
Maximum usable queues: 8, Queues in use: 5
    Scheduler map: <default>, Index: 2
    Congestion-notification: Disabled
Object      Name      Type      Index
Classifier   ipprec-compatibility  ip      13

Physical interface: ge-0/1/5, Index: 151
Maximum usable queues: 8, Queues in use: 5
    Scheduler map: <default>, Index: 2
    Congestion-notification: Disabled
Object      Name      Type      Index
Classifier   ipprec-compatibility  ip      13

Physical interface: ge-0/1/6, Index: 152
Maximum usable queues: 8, Queues in use: 5
    Scheduler map: <default>, Index: 2
    Congestion-notification: Disabled
Object      Name      Type      Index
Classifier   ipprec-compatibility  ip      13

Physical interface: ge-0/1/7, Index: 153
Maximum usable queues: 8, Queues in use: 5
    Scheduler map: <default>, Index: 2
    Congestion-notification: Disabled
Object      Name      Type      Index
Classifier   dl      dscp      11331

Physical interface: ge-0/2/0, Index: 154
Maximum usable queues: 8, Queues in use: 5
    Scheduler map: <default>, Index: 2
    Congestion-notification: Disabled
Object      Name      Type      Index
Classifier   ipprec-compatibility  ip      13

Physical interface: ge-0/2/1, Index: 155
Maximum usable queues: 8, Queues in use: 5
    Scheduler map: <default>, Index: 2
    Congestion-notification: Disabled
Object      Name      Type      Index

```

```

Classifier                ipprec-compatibility  ip                13

  Logical interface: ge-0/2/1.0, Index: 78

  Logical interface: ge-0/2/1.32767, Index: 79

Physical interface: xe-0/3/0, Index: 156
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object      Name                Type                Index
Classifier   ipprec-compatibility  ip                13

  Logical interface: xe-0/3/0.0, Index: 80

Physical interface: xe-0/3/1, Index: 157
Maximum usable queues: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object      Name                Type                Index
Classifier   ipprec-compatibility  ip                13

  Logical interface: xe-0/3/1.0, Index: 81

[edit]
user@host-g11#

```

show class-of-service interface (PPPoE Subscriber Interface for Enhanced Subscriber Management)

user@host> **show class-of-service interface pp0.3221225474**

```

  Logical interface: pp0.3221225475, Index: 3221225475
Object      Name                Type                Index
Traffic-control-profile TC_PROF_100_199_SERIES_UID1006 Output            4294967312
Scheduler-map      SMAP-1_UID1002      Output            4294967327
Rewrite-Output     ieee-rewrite        ieee8021p         60432
Rewrite-Output     rule1               ip                50463

  Adjusting application: PPPoE IA tags
    Adjustment type: absolute
    Configured shaping rate: 11000000
    Adjustment value: 5000000
    Adjustment target: node

```

```
Adjusting application: ucac  
Adjustment type: delta  
Configured shaping rate: 5000000  
Adjustment value: 100000  
Adjustment target: node
```


show class-of-service interface-set

Syntax

```
show class-of-service interface-set
<interface-set-name>
```

Release Information

Command introduced in Junos OS Release 9.4.
 Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description

Display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical class of service (CoS).

Options

- none**—Display CoS associations for all logical interface sets.
- interface-set *interface-set-name***—(Optional) Display CoS associations for the specified interface set.

Required Privilege Level

view

List of Sample Output

- [show class-of-service interface-set on page 1357](#)
- [show class-of-service interface-set \(Enhanced Subscriber Management\) on page 1357](#)

Output Fields

[Table 60 on page 1355](#) describes the output fields for the **show class-of-service interface-set** command. Output fields are listed in the approximate order in which they appear.

Table 60: show class-of-service interface-set Output Fields

Field Name	Field Description
Interface-set	Name of a logical interface set composed of one or more logical interfaces for which hierarchical scheduling is enabled.
Index	Index number of this interface set or the internal index number of this object.
Physical interface	Name of a physical interface.
Queues supported	Number of queues you can configure on the interface.

Table 60: show class-of-service interface-set Output Fields (*continued*)

Field Name	Field Description
Queues in use	Number of queues currently configured.
Output traffic control profile	Name of the output traffic control profile attached to the logical interface set.
Output traffic control profile remaining	(Enhanced subscriber management for MX Series routers) For dynamic subscriber management, name of the output traffic control profile for remaining traffic attached to the logical interface set.
Adjusting application	<p>Name of the application that communicates shaping-rate adjustment information to the Junos OS class-of-service process (cosd) on the broadband services router (BSR). The BSR uses the information from this application to perform shaping-rate adjustments on the scheduler node that manages the interface set. The adjusting application appears as ancp LS-0 which is the Junos OS Access Node Control Profile process (ancpd) that performs shaping-rate adjustments on schedule nodes. The nodes are logical interface sets configured to represent subscriber local loops. When the synchronization speed of the DSL line changes, ancpd communicates the local loop speed to cosd over the default logical system, LS-0, and then the BSR throttles the shaping rate on the scheduler node to the loop speed.</p> <p>The adjusting application can also appear as PPPoE, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual data rate downstream attribute. The overhead accounting value is based on the access loop encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).</p>
Adjustment type	Type of shaping-rate adjustment performed by the BSR on the scheduler node. The type of adjustment appears as Adjustment type , meaning that the configured shaping rate is adjusted by an absolute value as opposed to by a percentage of the configured rate.
Configured shaping rate	The maximum transmission rate on the physical interface as configured by the output traffic-control profile attached to the scheduler node.
Adjustment value	Value of the shaping-rate adjustment information sent by the adjusting application to cosd .
Adjustment overhead-accounting mode	Configured shaping mode: frame or cell .

Sample Output

show class-of-service interface-set

```
user@host> show class-of-service interface-set example-ifset-ge-4/0/0-7
```

```
Interface-set: example-ifset-ge-4/0/0-7, Index: 8
Physical interface: ge-4/0/0, Index: 270
Queues supported: 8, Queues in use: 8
  Output traffic control profile: example-tcp-basic-rate, Index: 11395
Adjusting application: ancp LS-0
  Adjustment type: absolute
  Configured shaping rate: 50000000
  Adjustment value: 888000
  Adjustment overhead-accounting mode: cell
```

show class-of-service interface-set (Enhanced Subscriber Management)

```
user@host> show class of service interface-set
```

```
Interface-set: ge-1/0/0-201-201, Index: 1
Physical interface: ge-1/0/0, Index: 142
Queues supported: 8, Queues in use: 4
  Output traffic control profile: LEVEL_2_UID1001, Index: 4294967307
  Output traffic control profile remaining: TCP_REMAIN_UID1003, Index: 4294967308
```

show class-of-service scheduler-hierarchy interface

Syntax

```
show class-of-service scheduler-hierarchy interface interface-name <detail>
```

Release Information

Command introduced in Junos OS Release 13.3 for MX Series Routers.
 Support for up to four hierarchy levels added in Junos OS Release 16.1.

NOTE: Before Junos OS R19.2, the shaping rate would incorrectly display as 90% of the guaranteed rate.

Description

For MPC/MIC interfaces only, display the scheduler hierarchy as well as the shaping rate, guaranteed rate, priorities, and queue weight information for each forwarding class at each hierarchy level.

Options

detail—(Optional) Display scheduler hierarchies based on the interface set.

interface-name—Display information about a specific interface.

Required Privilege Level

view

RELATED DOCUMENTATION

[hierarchical-scheduler \(Subscriber Interfaces on MX Series Routers\)](#) | 914

List of Sample Output

[show class-of-service scheduler-hierarchy interface on page 1359](#)

Output Fields

[Table 61 on page 1358](#) describes the output fields for the **show class-of-service scheduler-hierarchy interface** command. Output fields are listed in the approximate order in which they appear.

Table 61: show class-of-service scheduler-hierarchy interface Output Fields

Field Name	Field Description
interface	Interface name

Table 61: show class-of-service scheduler-hierarchy interface Output Fields (*continued*)

Field Name	Field Description
resource	Traffic resource associated with the logical interface
shaping-rate	Shaping rate in bits per second
guaranteed rate	Guaranteed rate in bits per second
guaranteed priority	Queue priority in the guaranteed region (high, low, or none)
excess priority	Queue priority in the excess region (high, low, or none)
queue weight	Queue weight for excess CoS weighted round-robin
excess weight	Interface unit per priority weights for excess weighted round-robin

Sample Output

show class-of-service scheduler-hierarchy interface

user@host> show class-of-service scheduler-hierarchy interface xe-1/0/0

Interface/ resource name	shaping rate kbits	guaranteed rate kbits	guaranteed/ excess priority	queue weight	excess weight high/low
xe-1/0/0	12000				
<<< L1					
xe-1/0/0 RTP	12000	0			1 1
best-effort	12000	0	Low Low	950	
network-control	12000	0	Low Low	50	
ifset1	12000	0			500 500
<<< L2					
ifset1 RTP	12000	0			1 1
be1	720	0	Low Low	250	
ncl	12000	0	Low Low	250	
demux0.96	3000	0			1 1
<<< L3					
demux0.96 RTP	3000	0			500 500

	bel	1000	0	Low	Low	250		
	ncl	3000	0	Low	Low	250		
	pp0.81	2000	0				1	1
<<<	L4							
	bel	1000	0	Low	Low	250		
	ncl	2000	0	Low	Low	250		

show class-of-service scheduler-hierarchy interface-set

Syntax

```
show class-of-service scheduler-hierarchy interface-set interface-set-name <detail>
```

Release Information

Command introduced in Junos OS Release 13.3 for MX Series Routers.

Description

For MPC/MIC interface sets only, display the scheduler hierarchy.

Options

detail—(Optional) Display scheduler hierarchies based on the interface-set.

interface-set-name—Display information about a specific interface-set.

Required Privilege Level

view

RELATED DOCUMENTATION

| [show interfaces queue](#)

List of Sample Output

[show class-of-service scheduler-hierarchy interface-set on page 1362](#)

Output Fields

[Table 62 on page 1361](#) describes the output fields for the **show class-of-service scheduler-hierarchy interface-set** command. Output fields are listed in the approximate order in which they appear.

Table 62: show class-of-service scheduler-hierarchy interface-set Output Fields

Field Name	Field Description
interface	Type of interface
resource	Traffic resource associated with the logical interface
shaping-rate	Actual shaping rate in bits per second
guaranteed rate	Actual guaranteed rate in bits per second
guaranteed priority	Actual queue priority in the guaranteed region (high, low, or none)

Table 62: show class-of-service scheduler-hierarchy interface-set Output Fields (*continued*)

Field Name	Field Description
excess priority	Actual queue priority in the excess region (high, low, or none)
queue weight	Actual queue weight for excess CoS weighted round-robin
excess weight	Actual interface-set per priority weights for excess weighted round-robin

Sample Output

show class-of-service scheduler-hierarchy interface-set

user@host> show class-of-service scheduler-hierarchy interface-set ifset

Interface/ resource name	shaping rate kbits	guaranteed rate kbits	guaranteed/ excess priority	queue weight	excess weight high/low
ge-1/0/0	100000				
ge-1/0/0 RTP	100000	0			1 1
be	100000	1000	Low Low	1	
da	9000	2000	Medium High	1	
vi	100000	3000	Medium None	626	
vo	100000	4000	High High	373	
gt	100000	0	High High	1	
ge-1/0/0.20	50000	40000			750 750
be	50000	1000	Low Low	1	
da	9000	2000	Medium High	1	
vi	50000	3000	Medium None	626	
vo	50000	4000	High High	373	
gt	50000	Disabled	High High	1	

show class-of-service scheduler-map

Syntax

```
show class-of-service scheduler-map  
<name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description

Display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.

Options

none—Display all scheduler maps.

name—(Optional) Display a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.

Required Privilege Level

view

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

List of Sample Output

[show class-of-service scheduler-map on page 1365](#)

[show class-of-service scheduler-map \(QFX Series\) on page 1366](#)

Output Fields

[Table 63 on page 1364](#) describes the output fields for the **show class-of-service scheduler-map** command. Output fields are listed in the approximate order in which they appear.

Table 63: show class-of-service scheduler-map Output Fields

Field Name	Field Description
Scheduler map	<p>Name of the scheduler map.</p> <p>(Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, SMAP-1_UID1002) instead of with a subscriber interface.</p>
Index	<p>Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles.</p> <p>(Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management.</p>
Scheduler	Name of the scheduler.
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
Transmit rate	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword remainder , which indicates that the scheduler receives the remaining bandwidth of the interface.
Rate Limit	Rate limiting configuration of the queue. Possible values are none , meaning no rate limiting, and exact , meaning the queue only transmits at the configured rate.
Maximum buffer delay	Amount of transmit delay (in milliseconds) or the buffer size of the queue. The buffer size is shown as a percentage of the total interface buffer allocation, or by the keyword remainder to indicate that the buffer is sized according to what remains after other scheduler buffer allocations.
Priority	Scheduling priority: low or high .
Excess priority	Priority of excess bandwidth: low , medium-low , medium-high , high , or none .
Explicit Congestion Notification	<p>(QFX Series, OCX Series, and EX4600 switches only) Explicit congestion notification (ECN) state:</p> <ul style="list-style-type: none"> • Disable—ECN is disabled on the specified scheduler • Enable—ECN is enabled on the specified scheduler <p>ECN is disabled by default.</p>
Adjust minimum	Minimum shaping rate for an adjusted queue, in bps.

Table 63: show class-of-service scheduler-map Output Fields (*continued*)

Field Name	Field Description
Adjust percent	Bandwidth adjustment applied to a queue, in percent.
Drop profiles	Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.
Loss priority	Packet loss priority for drop profile assignment.
Protocol	Transport protocol for drop profile assignment.
Name	Name of the drop profile.

Sample Output

show class-of-service scheduler-map

```
user@host> show class-of-service scheduler-map
```

```
Scheduler map: dd-scheduler-map, Index: 84
```

```
Scheduler: aa-scheduler, Index: 8721, Forwarding class: aa-forwarding-class
Transmit rate: 30 percent, Rate Limit: none, Maximum buffer delay: 39 ms,
Priority: high
```

```
Drop profiles:
```

Loss priority	Protocol	Index	Name
Low	non-TCP	8724	aa-drop-profile
Low	TCP	9874	bb-drop-profile
High	non-TCP	8833	cc-drop-profile
High	TCP	8484	dd-drop-profile

```
Scheduler: bb-scheduler, Forwarding class: aa-forwarding-class
Transmit rate: 40 percent, Rate limit: none, Maximum buffer delay: 68 ms,
Priority: high
```

```
Drop profiles:
```

Loss priority	Protocol	Index	Name
Low	non-TCP	8724	aa-drop-profile
Low	TCP	9874	bb-drop-profile
High	non-TCP	8833	cc-drop-profile
High	TCP	8484	dd-drop-profile

show class-of-service scheduler-map (QFX Series)

```
user@switch# show class-of-service scheduler-map
```

```
Scheduler map: be-map, Index: 12240
```

```
Scheduler:be-sched, Forwarding class: best-effort, Index: 115
```

```
Transmit rate: 30 percent, Rate Limit: none, Buffer size: remainder,
```

```
Buffer Limit: none, Priority: low
```

```
Excess Priority: unspecified, Explicit Congestion Notification: disable
```

```
Drop profiles:
```

Loss priority	Protocol	Index	Name
Low	any	3312	lan-dp
Medium-high	any	2714	be-dp1
High	any	3178	be-dp2

show class-of-service traffic-control-profile

Syntax

```
show class-of-service traffic-control-profile
<profile-name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Command introduced in Junos OS Release 12.2 for ACX Series Routers.

Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description

For Gigabit Ethernet IQ PICs, Channelized IQ PICs, EQ DPCs, and MPC/MIC interfaces only, display traffic shaping and scheduling profiles.

(ACX Series routers) For ATM IMA pseudowire interfaces, display traffic shaping and scheduling profiles.

Options

none—Display all profiles.

profile-name—(Optional) Display information about a single profile.

Required Privilege Level

view

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

List of Sample Output

[show class-of-service traffic-control-profile on page 1370](#)

[show class-of-service traffic-control-profile \(MX Series routers with Clear Channel Multi-Rate CE MIC\) on page 1370](#)

[show class-of-service traffic-control-profile \(ACX Series routers with ATM IMA pseudowire interfaces\) on page 1371](#)

[show class-of-service traffic-control-profile \(Enhanced Subscriber Management\) on page 1371](#)

Output Fields

[Table 64 on page 1368](#) describes the output fields for the **show class-of-service traffic-control-profile** command. Output fields are listed in the approximate order in which they appear.

Table 64: show class-of-service traffic-control-profile Output Fields

Field Name	Field Description
Traffic control profile	<p>Name of the traffic control profile.</p> <p>(Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, TC_PROF_100_199_SERIES_UID1000) instead of with a subscriber interface.</p>
Index	<p>Index number of the traffic control profile.</p> <p>(Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management.</p>
ATM Service	<p>(MX Series routers with ATM Multi-Rate CE MIC) Configured category of ATM service. Possible values:</p> <ul style="list-style-type: none"> • cbr—Constant bit rate. • rtvbr—Real time variable bit rate. • nrtvbr—Non real time variable bit rate. • ubr—Unspecified bit rate.
Maximum Burst Size	Configured maximum burst size, in cells.
Peak rate	Configured peak rate, in cps.
Sustained rate	Configured sustained rate, in cps.
Shaping rate	<p>Configured shaping rate, in bps.</p> <p>NOTE: (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.</p>
Shaping rate burst	<p>Configured burst size for the shaping rate, in bytes.</p> <p>NOTE: (MX Series routers with ATM Multi-Rate CE MIC) Configured maximum burst rate, in cells.</p>
Shaping rate priority high	Configured shaping rate for high-priority traffic, in bps.
Shaping rate priority medium	Configured shaping rate for medium-priority traffic, in bps.
Shaping rate priority low	Configured shaping rate for low-priority traffic, in bps.

Table 64: show class-of-service traffic-control-profile Output Fields (continued)

Field Name	Field Description
Shaping rate excess high	Configured shaping rate for high-priority excess traffic, in bps.
Shaping rate excess low	Configured shaping rate for low-priority excess traffic, in bps.
Scheduler map	<p>Name of the associated scheduler map.</p> <p>(Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, SMAP-1_UID1002) instead of with a subscriber interface.</p>
Delay Buffer rate	Configured delay buffer rate, in bps.
Excess rate	Configured excess rate, in percent or proportion.
Excess rate high	Configured excess rate for high priority traffic, in percent or proportion.
Excess rate low	Configured excess rate for low priority traffic, in percent or proportion.
Guaranteed rate	<p>Configured guaranteed rate, in bps or cps.</p> <p>NOTE: (MX Series routers with ATM Multi-Rate CE MIC) This value depends on the ATM service category chosen. Possible values:</p> <ul style="list-style-type: none"> • cbr—Guaranteed rate is equal to the configured peak rate in cps. • rtvbr—Guaranteed rate is equal to the configured sustained rate in cps. • nrtvbr—Guaranteed rate is equal to the configured sustained rate in cps.
Guaranteed rate burst	Configured burst size for the guaranteed rate, in bytes.
adjust-minimum	Configured minimum shaping rate for an adjusted queue, in bps.
overhead accounting mode	Configured shaping mode: Frame Mode or Cell Mode .
Overhead bytes	Configured byte adjustment value.
Adjust parent	<p>Configured shaping-rate adjustment for parent scheduler nodes. If enabled, this field appears.</p> <p>flow-aware indicates that the parent scheduler node is adjusted only once per multicast channel.</p>

Sample Output

show class-of-service traffic-control-profile

user@host> show class-of-service traffic-control-profile

```
Traffic control profile: Profile1, Index: 57625
  Scheduler map: m1
  Delay Buffer rate: 500000
  Guaranteed rate: 1000000

Traffic control profile: Profile2, Index: 57624
  Scheduler map: m2
  Delay Buffer rate: 600000
  Guaranteed rate: 2000000

Traffic control profile: Profile3, Index: 57627
  Scheduler map: m3
  Delay Buffer rate: 800000
  Guaranteed rate: 3000000
  .Excess rate high: proportion 4

Traffic control profile: Profile4, Index: 57626
  Scheduler map: m4
  Delay Buffer rate: 750000
  Guaranteed rate: 4000000
  ..adjust-minimum 20000000

Traffic control profile: foo, Index: 57626
  Shaping rate: 100000000
  Scheduler map: <default>
  Overhead accounting mode: Frame Mode
  Frame mode overhead accounting bytes: -12
  Adjust parent: flow-aware
```

show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC)

user@host> show class-of-service traffic-control-profile

```
Traffic control profile: at-vbr1, Index: 11395
  ATM Service: RTVBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode
  Shaping rate: 1000 cps
```



```

Shaping rate burst: 500 cells
Delay Buffer rate: 2000 cps
Guaranteed rate: 1000 cps

Traffic control profile: foo, Index: 38286
ATM Service: UBR
Scheduler map: m3
overhead accounting mode: Frame Mode

```

show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces)

user@host> **show class-of-service traffic-control-profile**

```

Traffic control profile: foo, Index: 38286
ATM Service: RTVBR
Shaping rate: 2000 cps
Shaping rate burst: 200 cells
Scheduler map: <default>
Delay Buffer rate: 1000 cps
Guaranteed rate: 1700 cps

```

show class-of-service traffic-control-profile (Enhanced Subscriber Management)

user@host> **show class-of-service traffic-control-profile**

```

Traffic control profile: TC_PROF_100_199_SERIES_UID1000, Index: 4294967313
Shaping rate: 11000000
Shaping rate burst: 1 bytes
Scheduler map: SMAP-1_UID1002
Delay Buffer rate: 5000000
Overhead accounting mode: Cell Mode
Frame mode overhead accounting bytes: -4
Cell mode overhead accounting bytes: 20

```

show dynamic-profile session

Syntax

```
show dynamic-profile session
<client-id client-id>
<profile-name profile-name>
<service-id service-id>
```

Release Information

Command introduced in Junos OS Release 13.3.

Description

Display dynamic profile (client or service) information for all subscribers or for subscribers specified by client ID or service session ID. You can filter the output by also specifying a dynamic profile.

NOTE:

- The output does not display the variable stanzas defined in the dynamic profile configuration.
- The variables in the profile configuration are replaced with subscriber specific values.
- If the conditional variable in the dynamic profile is evaluated as NULL, the subscriber value for the variable is displayed as **NONE** in the command output.
- The variable is also displayed as **NONE** when the variable (any variable and not necessarily conditional) in the dynamic profile has no value associated with it.
- The format in which the configuration is displayed looks similar, but not exactly the same as the format of the **show configuration dynamic-profiles** command.

Options

client-id *client-id*—Display dynamic profile information for subscribers associated with the specified client.

profile-name *profile-name*—(Optional) Display dynamic profile information for the specified subscriber or service profile.

service-id *service-id*—Display dynamic profile information for subscribers associated with the specified service session.

Required Privilege Level

view

List of Sample Output

[show dynamic-profile session client-id \(Client ID\) on page 1373](#)

[show dynamic-profile session client-id profile-name \(Client ID and Dynamic Profile\) on page 1375](#)

[show dynamic-profile session service-id \(Service Session\) on page 1376](#)

Output Fields

This command displays the dynamic client or service profile configuration for each subscriber.

Sample Output

show dynamic-profile session client-id (Client ID)

user@host>show dynamic-profile session client-id 20

```
pppoe {
  interfaces {
    pp0 {
      unit 1073741831 {
        ppp-options {
          chap;
          pap;
        }
        pppoe-options {
          underlying-interface ge-2/0/0.0;
          server;
        }
        family {
          inet {
            unnumbered-address lo0.0;
          }
        }
      }
    }
  }
}
class-of-service {
  traffic-control-profiles {
    tcp1 {
      scheduler-map smap1_UID1024;
      shaping-rate 100m;
    }
  }
  interfaces {
    pp0 {
      unit 1073741831 {
        output-traffic-control-profile tcp1;
      }
    }
  }
}
```

```

    }
  }
}
scheduler-maps {
  smap1_UID1024 {
    forwarding-class best-effort scheduler sch1_UID1023;
  }
}
schedulers {
  sch1_UID1023 {
    transmit-rate percent 40;
    buffer-size percent 40;
    priority low;
  }
}
}
}
filter-service {
  interfaces {
    pp0 {
      unit 1073741831 {
        family {
          inet {
            filter {
              input input-filter_UID1026 precedence 50;
              output output-filter_UID1027 precedence 50;
            }
          }
        }
      }
    }
  }
}
}
firewall {
  family {
    inet {
      filter input-filter_UID1026 {
        interface-specific;
        term t1 {
          then {
            policer policer1_UID1025;
            service-accounting;
          }
        }
        term rest {

```



```

cos-service {
  class-of-service {
    scheduler-maps {
      smap2_UID1029 {
        forwarding-class assured-forwarding scheduler sch2_UID1028;
      }
    }
    schedulers {
      sch2_UID1028 {
        transmit-rate percent 60;
        buffer-size percent 60;
        priority high;
      }
    }
  }
}

```

show dynamic-profile session service-id (Service Session)

user@host>show dynamic-profile session service-id 21

```

filter-service {
  interfaces {
    pp0 {
      unit 1073741831 {
        family {
          inet {
            filter {
              input input-filter_UID1026 precedence 50;
              output output-filter_UID1027 precedence 50;
            }
          }
        }
      }
    }
  }
}

firewall {
  family {
    inet {
      filter input-filter_UID1026 {
        interface-specific;
        term t1 {
          then {

```

```
        policer policer1_UID1025;  
        service-accounting;  
    }  
}  
term rest {  
    then accept;  
}  
}  
filter output-filter_UID1027 {  
    interface-specific;  
    term rest {  
        then accept;  
    }  
}  
}  
}  
policer policer1_UID1025 {  
    if-exceeding {  
        bandwidth-limit 1m;  
        burst-size-limit 15k;  
    }  
    then discard;  
}  
}  
}
```

show firewall

List of Syntax

[Syntax on page 1378](#)

[Syntax \(EX Series Switches\) on page 1378](#)

Syntax

```
show firewall
<application (CFM | eswd | RMPS)>>
<counter counter-name>
<detail>
<filter filter-name>
<filter regex regular-expression>
<logical-system (all | logical-system-name)>
<terse>
```

Syntax (EX Series Switches)

```
show firewall
<application (CFM | eswd | RMPS)>>
<counter counter-name>
<detail>
<filter filter-name>
<filter regex regular-expression>
<log <(detail | interface interface-name)>>
<policer counters <(detail | counter-id counter-index <detail>)>>
<terse>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Option **logical-system** introduced in Junos OS Release 9.3.

Option **terse** introduced in Junos OS Release 9.4.

Option **policer counters** introduced in Junos OS Release 12.2 for EX Series switches.

Option **detail** introduced in Junos OS Release 12.3 for EX Series switches.

Option **detail** introduced in Junos OS Release 14.1 for MX Series routers.

Option **regex *regular-expression*** introduced in Junos OS Release 14.2.

Option **lsp** introduced in Junos OS Evolved Release 18.3R1.

Description

Display enhanced statistics and counters for all configured firewall filters.

If you query for options on the **show firewall filter** command, on Junos OS systems, you will see this output, which includes the configured Flowspec filters:

```
show firewall filter ?
```

```
Possible completions:
<filtername>          Filter name
__flowspec_default_inet__  # Flowspec filter name
application           Owner application
counter               Counter name
logical-system        Name of logical system, or 'all'
regex                 Show filter using regular expression
version               Show filter version installed
```

However, on Junos OS Evolved systems, the Flowspec filters names are not shown here. To view Flowspec filters, use the **show firewall application routing** command.

Options

none—(Optional) Display statistics and counters for all configured firewall filters and counters. For EX Series switches, this command also displays statistics about all configured policers.

application (CFM | eswd | RMPS)—(Optional) Show firewall elements owned by the selected software component:

- Connectivity Fault Management (CFM)
- Ethernet switching daemon (eswd)—Shows only on devices that support it.
- Resource Management and Packet Steering (RMPS)

counter counter-name—(Optional) Name of a filter counter.

detail—(EX Series switches and MX Series routers only) (Optional) Display firewall filter statistics and enhanced policer statistics and counters.

filter filter-name—(Optional) Name of a configured filter.

filter regex regular-expression—(Optional) Regular expression that matches the names of a subset of filters.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.

log—(Optional) Display log entries for firewall filters.

log <(detail | interface interface-name)>—(EX Series switches only) (Optional) Display detailed log entries of firewall activity or log information about a specific interface.

policer counters <(detail | counter-id *counter-index* <detail>)>—(EX8200 switches only) (Optional) Display enhanced policer counter statistics in brief or in detail.

terse—(Optional) Display firewall filter names only.

Required Privilege Level

view

RELATED DOCUMENTATION

clear firewall 1256
show firewall log 1389
<i>Verifying That Firewall Filters Are Operational</i>
<i>Verifying That Policers Are Operational</i>
<i>show policer</i>
Enhanced Policer Statistics Overview 390
enhanced-policer 846

List of Sample Output

- [show firewall filter \(MX Series Router and EX Series Switch\) on page 1383](#)
- [show firewall filter \(non MX Series Router and EX Series Switch\) on page 1383](#)
- [show firewall filter \(Dynamic Input Filter\) on page 1384](#)
- [show firewall \(Logical Systems\) on page 1384](#)
- [show firewall \(counter counter-name\) on page 1384](#)
- [show firewall log on page 1385](#)
- [show firewall policer counters \(EX8200 Switch\) on page 1385](#)
- [show firewall policer counters \(detail\) \(EX8200 Switch\) on page 1386](#)
- [show firewall policer counters \(counter-id counter-index\) \(EX8200 Switch\) on page 1386](#)
- [show firewall policer counters \(counter-id counter-index detail\) \(EX8200 Switch\) on page 1387](#)
- [show firewall detail on page 1387](#)
- [show firewall application cfm \(Junos OS Evolved\) on page 1387](#)

Output Fields

[Table 65 on page 1381](#) lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

Table 65: show firewall Output Fields

Field Name	Field Description
Filter	<p>Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.</p> <p>Except on EX Series switches:</p> <ul style="list-style-type: none"> • When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either -i for an input filter or -o for an output filter. • When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either -in for an input filter or -out for an output filter. When a logical system-specific filter is displayed, the name of the filter is prefixed with two underscore (_) characters and the name of the logical system (for example, _ls1/filter1). • When a service filter is displayed that uses a service set, the separator between the service-set name and the service-filter name is a semicolon (;). <p>NOTE: For bridge family filter, the ip-protocol match criteria is supported only for IPv4 and not for IPv6. This is applicable for line cards that support the Junos Trio chipset, such as the MX 3D MPC line cards.</p>
Counters	<p>Display filter counter information:</p> <ul style="list-style-type: none"> • Name—Name of a filter counter that has been configured with the counter firewall filter action. • Bytes—Number of bytes that match the filter term under which the counter action is specified. • Packets—Number of packets that matched the filter term under which the counter action is specified. <p>NOTE: On M and T Series routers, firewall filters cannot count ip-options packets on a per option type and per interface basis. A limited work around is to use the show pfe statistics ip options command to see ip-options statistics on a per Packet Forwarding Engine (PFE) basis. See <i>show pfe statistics ip</i> for sample output.</p>

Table 65: show firewall Output Fields (*continued*)

Field Name	Field Description
Policers	<p>Display policer information:</p> <ul style="list-style-type: none"> • Name—Name of policer. • Bytes—(For two-color policers on MX Series routers and EX Series switches, and for hierarchical policers on interfaces hosted on MICs and MPCs in MX Series routers) Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. For other combinations of policer type, device, and line card type, this field is blank. • Packets—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.
Policer Counter Index	(EX8200 switch only) Global management counter ID. The counter ID value (<i>counter-index</i>) can be 0, 1, or 2.
Green	(EX8200 switch only) Number of packets within the limits. The number of packets is smaller than the committed information rate (CIR).
Yellow	(EX8200 switch only) Number of packets partially within the limits. The number of packets is greater than the CIR, but the burst size is within the excess burst size (EBS) limit.
Discard	(EX8200 switch only) Number of discarded packets.
Bytes	(EX8200 switch only) Number of green, yellow, red, or discarded packets in bytes.
Packets	(EX8200 switch only) Number of green, yellow, red, or discarded packets.
Filter name	(EX8200 switch only) Name of the filter with a term associated to a policer.
Term name	(EX8200 switch only) Name of the term associated with a policer.
Policer name	(EX8200 switch only) Name of the policer that is associated with a global management counter.

Table 65: show firewall Output Fields (*continued*)

Field Name	Field Description
P1-t1	<ul style="list-style-type: none"> • 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 in this counter. • Offered packet statistics for traffic subjected to policing. • 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.

Sample Output

show firewall filter (MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
```

```
Filter: test
Counters:
Name                Bytes          Packets
Counter-1           0              0
Counter-2           0              0
Policers:
Name                Bytes          Packets
Policer-1          2770           70
```

show firewall filter (non MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
```

```
Filter: test
Counters:
Name                Bytes          Packets
Counter-1           0              0
Counter-2           0              0
Policers:
Name                Bytes          Packets
Policer-1           70
```

show firewall filter (Dynamic Input Filter)

```
user@host> show firewall filter dfwd-ge-5/0/0.1-in
```

```
Filter: dfwd-ge-5/0/0.1-in
Counters:
Name                                     Bytes      Packets
cl-ge-5/0/0.1-in                        0          0
```

show firewall (Logical Systems)

```
user@host> show firewall
```

```
Filter: __lr1/test
Counters:
Name                                     Bytes      Packets
icmp                                     420        5
Filter: __default_bpdu_filter__
Filter: __lr1/inet_filter1
Counters:
Name                                     Bytes      Packets
inet_tcp_count                          0          0
inet_udp_count                          0          0
Filter: __lr1/inet_filter2
Counters:
Name                                     Bytes      Packets
inet_icmp_count                         0          0
inet_pim_count                          0          0
Filter: __lr2/inet_filter1
Counters:
Name                                     Bytes      Packets
inet_tcp_count                          0          0
inet_udp_count                          0          0
```

show firewall (counter counter-name)

```
user@host> show firewall counter icmp-counter
```

```
Filter: ingress-port-voip-class-filter
Counters:
```

Name	Bytes	Packets
icmp-counter	0	0

show firewall log

user@host> show firewall log

Log :

Time	Filter	Action	Interface	Protocol	Src Addr
	Dest Addr				
08:00:53	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:52	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:51	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:50	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:49	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:48	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				
08:00:47	pfe	R	ge-1/0/1.0	ICMP	192.168.3.5
	192.168.3.4				

show firewall policer counters (EX8200 Switch)

user@switch> show firewall policer counters

Policer Counter Index 0:

	Bytes	Packets
Green:	73	15914
Yellow:	9	1962
Discard:	119	25942

Policer Counter Index 1:

	Bytes	Packets
Green:	0	0
Yellow:	0	0
Discard:	0	0

Policer Counter Index 2:

	Bytes	Packets
Green:	0	0
Yellow:	0	0
Discard:	0	0

show firewall policer counters (detail) (EX8200 Switch)

```
user@switch> show firewall policer counters detail
```

Policer Counter Index 0:

	Bytes	Packets
Green:	73	15914
Yellow:	9	1962
Discard:	119	25942

Filter name	Term name	Policer name
myfilter	polcr-term-1	myfilter-polcr-1
inet-filter-ae	ae-snmp	policer-1
inet-filter-ae	ae-ssh	policer-2

Policer Counter Index 1:

	Bytes	Packets
Green:	0	0
Yellow:	0	0
Discard:	0	0

Filter name	Term name	Policer name
-------------	-----------	--------------

Policer Counter Index 2:

	Bytes	Packets
Green:	0	0
Yellow:	0	0
Discard:	0	0

Filter name	Term name	Policer name
-------------	-----------	--------------

show firewall policer counters (counter-id counter-index) (EX8200 Switch)

```
user@switch> show firewall policer counters counter-id 0
```

Policer Counter Index 0:

Bytes	Packets
-------	---------

Green:	73	15914
Yellow:	9	1962
Discard:	119	25942

show firewall policer counters (counter-id counter-index detail) (EX8200 Switch)

user@switch> show firewall policer counters counter-id 0 detail

```

Policer Counter Index 0:

              Bytes           Packets
Green:         73             15914
Yellow:        9              1962
Discard:       119           25942

Filter name    Term name      Policer name
myfilter       polcr-term-1    myfilter-polcr-1
inet-filter-ae ae-snmpp             policer-1
inet-filter-ae ae-ssh              policer-2

```

show firewall detail

user@host> show firewall detail

```

Filter: __default_bpdu_filter__

Filter: foo
Counters:
Name              Bytes           Packets
cl                17652140         160474
Policers:
Name              Bytes           Packets
Pl-tl
  OOS              0              18286
  Offered          0 18446744073709376546
  Transmitted      0 18446744073709358260

```

show firewall application cfm (Junos OS Evolved)

user@host> show firewall application cfm

Filter: __cfm_filter_et-0/0/0__

Counters:

Name	Bytes	Packets
__cfm_cc_term_lvl_0__	0	0
__cfm_cc_term_lvl_1__	0	0
__cfm_cc_term_lvl_2__	0	0
__cfm_cc_term_lvl_3__	0	0
__cfm_cc_term_lvl_4__	0	0
__cfm_cc_term_lvl_5__	0	0
__cfm_cc_term_lvl_6__	0	0
__cfm_cc_term_lvl_7__	0	0
__cfm_ethtype_term__	0	0
__cfm_lt_term_lvl_0__	0	0
__cfm_lt_term_lvl_1__	0	0
__cfm_lt_term_lvl_2__	0	0
__cfm_lt_term_lvl_3__	0	0
__cfm_lt_term_lvl_4__	0	0
__cfm_lt_term_lvl_5__	0	0
__cfm_lt_term_lvl_6__	0	0
__cfm_lt_term_lvl_7__	0	0
__cfm_ucast_term_536__	0	0

show firewall log

List of Syntax

[Syntax on page 1389](#)

[Syntax \(EX Series Switches\) on page 1389](#)

Syntax

```
show firewall log
<detail>
<extensive>
<interface interface-name>
<logical-system (logical-system-name | all)>
```

Syntax (EX Series Switches)

```
show firewall log
<detail>
<interface interface-name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

extensive option introduced in Junos OS Release 16.1.

logical-system option introduced in Junos OS Release 9.3.

Description

Display log information about firewall filters.

Options

none—Display log information about firewall filters.

detail—(Optional) Display detailed information.

extensive—(Optional) Display hex dump of packet captured by log action.

interface *interface-name*—(Optional) Display log information about a specific interface.

logical-system (*logical-system-name* | all)—(Optional) Perform this operation on all logical systems or on a particular system.

Required Privilege Level

view

List of Sample Output

[show firewall log on page 1390](#)

[show firewall log detail on page 1391](#)

[show firewall log extensive on page 1391](#)

Output Fields

Table 66 on page 1390 lists the output fields for the **show firewall log** command. Output fields are listed in the approximate order in which they appear.

Table 66: show firewall log Output Fields

Field Name	Field Description
Time of Log	Time that the event occurred.
Filter	<ul style="list-style-type: none"> Displays the name of a configured firewall filter or service filter only if the packet hit the filter's log action in a kernel filter (in the control plane). For any traffic that reaches the Routing Engine, the packets hit the log action in the kernel. For all other logged packets (packet hit the filter's log action in the Packet Forwarding Engine), this field displays pfe instead of a configured filter name.
Filter Action	Filter action: <ul style="list-style-type: none"> A—Accept D—Discard R—Reject
Name of Interface	<ul style="list-style-type: none"> Displays a physical interface name if the packet arrived at a port on a line card. Displays local if the packet was generated by the device's internal Ethernet interface, em1 or fxp1, which connects the Routing Engine with the router's packet-forwarding components.
Name of protocol	Packet's protocol name: egp , gre , icmp , ipip , ospf , pim , rsvp , tcp , or udp .
Packet length	Length of the packet.
Source address	Packet's source address.
Destination address	Packet's destination address and port.

Sample Output

show firewall log

```
user@host>show firewall log
```

Time	Filter	Action	Interface	Protocol	Src Addr	Dest Addr
13:10:12	pfe	D	rlsq0.902	ICMP	192.0.2.2	192.0.2.1
13:10:11	pfe	D	rlsq0.902	ICMP	192.0.2.2	192.0.2.1

show firewall log detail

user@host> show firewall log detail

```
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
....
```

show firewall log extensive

user@host> show firewall log extensive

```
Time of Log: 2016-01-17 22:16:21 PST, Filter: pfe, Filter action: accept, Name of
interface: xe-0/0/1.0
Name of protocol: UDP, Packet Length: 98, Source address: 203.0.113.1, Destination
address: 203.0.113.1
```

```
: 00-0F: 00 01 03 ee ee ff 00 01 - 09 22 55 ee 81 00 02 58
: 10-1F: 08 00 45 00 00 62 00 00 - 00 00 40 11 77 8a 01 00
: 20-2F: 00 01 02 00 00 01 1c 00 - 1c 00 00 4e 19 83 00 01
: 30-3F: 02 03 04 05 06 07 08 09 - 0a 0b 0c 0d 0e 0f 10 11
: 40-4F: 12 13 14 15 16 17 18 19 - 1a 1b 1c 1d 1e 1f 20 21
: 50-5F: 22 23 24 25 26 27 28 29 - 2a 2b 00 00 00 00 00 00
: 60-6F: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00
: 70-7F: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00
```

show firewall templates-in-use

Syntax

```
show firewall templates-in-use
```

Release Information

Command introduced in Junos OS Release 12.3.

Description

Display the names of configured filter templates that are currently in use by dynamic subscribers and the number of times each template is referenced.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear firewall](#) | [1256](#)

[show firewall log](#) | [1389](#)

List of Sample Output

[show firewall templates-in-use on page 1394](#)

Output Fields

[Table 67 on page 1393](#) lists the output fields for the **show firewall templates-in-use** command. Output fields are listed in the approximate order in which they appear.

Table 67: show firewall templates-in-use Output Fields

Field Name	Field Description
Filter Template	Name of a filter that has been configured using the filter statement at either the [edit firewall] or [edit dynamic-profiles <i>profile-name</i> firewall] hierarchy and is being used as a template for dynamic subscriber filtering.
Reference Count	Number of times the filter has been referenced by subscribers accessing the network.

Sample Output

show firewall templates-in-use

user@host> show firewall templates-in-use

Dynamic Subscribers Reference Counts	
Filter Template	Reference Count
-----	-----
egressFilter	10
ingressFilter	10
dfilter	5
dfilter-pol	5

show igmp group

List of Syntax

[Syntax on page 1395](#)

[Syntax \(EX Series Switch and the QFX Series\) on page 1395](#)

Syntax

```
show igmp group
<brief | detail>
<group-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
show igmp group
<brief | detail>
<group-name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.3 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display Internet Group Management Protocol (IGMP) group membership information.

Options

none—Display standard information about membership for all IGMP groups.

brief | detail—(Optional) Display the specified level of output.

group-name—(Optional) Display group membership for the specified IP address only.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear igmp membership](#) | [1259](#)

List of Sample Output

[show igmp group \(Include Mode\) on page 1397](#)

[show igmp group \(Exclude Mode\) on page 1397](#)

[show igmp group brief on page 1398](#)

[show igmp group detail on page 1398](#)

Output Fields

[Table 68 on page 1396](#) describes the output fields for the **show igmp group** command. Output fields are listed in the approximate order in which they appear.

Table 68: show igmp group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the IGMP membership report. A name of local indicates that the local routing device joined the group itself.	All levels
Group	Group address.	All levels
Group Mode	Mode the SSM group is operating in: Include or Exclude .	All levels
Source	Source address.	All levels
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	detail
Last reported by	Address of the host that last reported membership in this group.	All levels
Timeout	Time remaining until the group membership is removed.	brief none
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	detail
Type	Type of group membership: <ul style="list-style-type: none"> • Dynamic—Host reported the membership. • Static—Membership is configured. 	All levels

Sample Output

show igmp group (Include Mode)

user@host> show igmp group

```
Interface: t1-0/1/0.0
  Group: 198.51.100.1
    Group mode: Include
    Source: 203.0.113.2
    Last reported by: 203.0.113.52
    Timeout:      24 Type: Dynamic
  Group: 198.51.100.1
    Group mode: Include
    Source: 203.0.113.3
    Last reported by: 203.0.113.52
    Timeout:      24 Type: Dynamic
  Group: 198.51.100.1
    Group mode: Include
    Source: 203.0.113.4
    Last reported by: 203.0.113.52
    Timeout:      24 Type: Dynamic
  Group: 198.51.100.2
    Group mode: Include
    Source: 203.0.113.4
    Last reported by: 203.0.113.52
    Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 198.51.100.12
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic
  Group: 198.51.100.22
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic
```

show igmp group (Exclude Mode)

user@host> show igmp group

```

Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 198.51.100.2
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic
  Group: 198.51.100.22
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic

```

show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

show igmp group detail

user@host> **show igmp group detail**

```

Interface: t1-0/1/0.0
  Group: 198.51.100.1
    Group mode: Include
    Source: 203.0.113.2
    Source timeout: 12
    Last reported by: 203.0.113.52
    Group timeout:      0 Type: Dynamic
  Group: 198.51.100.1
    Group mode: Include
    Source: 203.0.113.3
    Source timeout: 12
    Last reported by: 203.0.113.52
    Group timeout:      0 Type: Dynamic
  Group: 198.51.100.1
    Group mode: Include
    Source: 203.0.113.4
    Source timeout: 12
    Last reported by: 203.0.113.52
    Group timeout:      0 Type: Dynamic
  Group: 198.51.100.2
    Group mode: Include
    Source: 203.0.113.4

```

```
Source timeout: 12
Last reported by: 203.0.113.52
Group timeout:      0 Type: Dynamic
Interface: tl-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
Group: 198.51.100.12
Group mode: Exclude
Source: 0.0.0.0
Source timeout: 0
Last reported by: Local
Group timeout:      0 Type: Dynamic
Group: 198.51.100.22
Group mode: Exclude
Source: 0.0.0.0
Source timeout: 0
Last reported by: Local
Group timeout:      0 Type: Dynamic
```

show igmp interface

List of Syntax

[Syntax on page 1400](#)

[Syntax \(EX Series Switches and the QFX Series\) on page 1400](#)

Syntax

```
show igmp interface
<brief | detail>
<interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and the QFX Series)

```
show igmp interface
<brief | detail>
<interface-name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.3 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.

Options

none—Display standard information about all IGMP-enabled interfaces.

brief | detail—(Optional) Display the specified level of output.

interface-name—(Optional) Display information about the specified IGMP-enabled interface only.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear igmp membership](#) | [1259](#)

List of Sample Output

[show igmp interface on page 1403](#)

[show igmp interface brief on page 1404](#)

[show igmp interface detail on page 1404](#)

[show igmp interface <interface-name> on page 1404](#)

Output Fields

[Table 69 on page 1401](#) describes the output fields for the **show igmp interface** command. Output fields are listed in the approximate order in which they appear.

Table 69: show igmp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the routing device that has been elected to send membership queries.	All levels
State	State of the interface: Up or Down .	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.	All levels
Timeout	How long until the IGMP querier is declared to be unreachable, in seconds.	All levels
Version	IGMP version being used on the interface: 1 , 2 , or 3 .	All levels
Groups	Number of groups on the interface.	All levels
Group limit	Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected.	All levels
Group threshold	Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.	All levels
Group log-interval	Time (in seconds) between consecutive log messages.	All levels

Table 69: show igmp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Immediate Leave	<p>State of the immediate leave option:</p> <ul style="list-style-type: none"> • On—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface. • Off—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds. 	All levels
Promiscuous Mode	<p>State of the promiscuous mode option:</p> <ul style="list-style-type: none"> • On—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces. • Off—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces. 	All levels
Distributed	<p>State of IGMP, which, by default, takes place on the Routing Engine for MX Series routers but can be distributed to the Packet Forwarding Engine to provide faster processing of join and leave events.</p> <ul style="list-style-type: none"> • On—distributed IGMP is enabled. 	All levels
Passive	<p>State of the passive mode option:</p> <ul style="list-style-type: none"> • On—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves. • Off—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves. <p>The passive statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the on state declaration:</p> <ul style="list-style-type: none"> • send-general-query—The interface sends general queries. • send-group-query—The interface sends group-specific and group-source-specific queries. • allow-receive—The interface receives control traffic. 	All levels
OIF map	Name of the OIF map (if configured) associated with the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map (if configured) used on the interface.	All levels

Table 69: show igmp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Configured Parameters	<p>Information configured by the user:</p> <ul style="list-style-type: none"> • IGMP Query Interval—Interval (in seconds) at which this router sends membership queries when it is the querier. • IGMP Query Response Interval—Time (in seconds) that the router waits for a report in response to a general query. • IGMP Last Member Query Interval—Time (in seconds) that the router waits for a report in response to a group-specific query. • IGMP Robustness Count—Number of times the router retries a query. 	All levels
Derived Parameters	<p>Derived information:</p> <ul style="list-style-type: none"> • IGMP Membership Timeout—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. • IGMP Other Querier Present Timeout—Time (in seconds) that the router waits for the IGMP querier to send a query. 	All levels

Sample Output

show igmp interface

```
user@host> show igmp interface
```

```

Interface: at-0/3/1.0
  Querier: 203.0.3.113.31
  State:          Up Timeout:    None Version:  2 Groups:      4
  SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
  Querier: 203.0.113.11
  State:          Up Timeout:    None Version:  2 Groups:      2
  SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
  Querier: 203.0.113.21
  State:          Up Timeout:    None Version:  2 Groups:      4
  SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off
Passive: Off

```

```
Distributed: OnConfigured Parameters:
```

```
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

```
Derived Parameters:
```

```
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 1403](#).

show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 1403](#).

show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
```

```
Interface: ge-3/2/0.0
  Querier: 203.0.113.111
  State: Up Timeout:    None
  Version: 3
  Groups: 1
  Group limit: 8
  Group threshold: 60
  Group log-interval: 10
  Immediate leave: Off
  Promiscuous mode: Off
  Distributed: On
```

show interfaces statistics

Syntax

```
show interfaces statistics interface-name
<satellite-device [device-alias-name |all ]>
<detail>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Command introduced in Junos OS Release 12.2 for ACX Series Routers.

satellite-device option introduced in Junos OS Release 14.2R3.

Description

Display static interface statistics, such as errors.

NOTE: When the **show interfaces statistics** command is executed on an interface that is configured on T4000 Type 5 FPC, the *IPv6 transit statistics* field displays:

- Total statistics (sum of transit and local statistics) at the physical interface level
- Transit statistics at the logical interface level

Options

interface-name—Name of an interface.

satellite-device [*device-alias-name* | all]—(Junos Fusion only) (Optional) Display interface statistics for interfaces on the specified satellite device in the Junos Fusion, or on all satellite devices in the Junos Fusion.

NOTE: In a Junos Fusion Enterprise, logical interface statistics are not synced across aggregation devices in a dual-aggregation device topology.

detail—(Optional) Display detailed output.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear interfaces statistics](#) | [1263](#)

List of Sample Output

[show interfaces statistics \(Fast Ethernet\) on page 1406](#)
[show interfaces statistics \(Gigabit Ethernet PIC—Egress\) on page 1407](#)
[show interfaces statistics detail \(Aggregated Ethernet\) on page 1410](#)
[show interfaces statistics detail \(Aggregated Ethernet—Ingress\) on page 1411](#)
[show interfaces statistics detail \(Aggregated Ethernet—Egress\) on page 1413](#)
[show interfaces statistics \(SONET/SDH\) on page 1414](#)
[show interfaces statistics \(Aggregated SONET/SDH—Ingress\) on page 1416](#)
[show interfaces statistics \(Aggregated SONET/SDH—Egress\) on page 1417](#)
[show interfaces statistics \(MX Series Routers\) on page 1418](#)
[show interfaces statistics \(MX Series Routers: Dynamic Interfaces with RPF Check Detail\) on page 1419](#)
[show interfaces statistics \(PTX Series Packet Transport Routers\) on page 1420](#)
[show interfaces statistics \(ACX Series routers\) on page 1421](#)

Output Fields

Output from both the **show interfaces *interface-name* detail** and the **show interfaces *interface-name* extensive** commands include all the information displayed in the output from the **show interfaces statistics** command. For more information, see the particular interface type in which you are interested. For information about destination class and source class statistics, see the “Destination Class Field” section and the “Source Class Field” section under *Common Output Fields Description*. For information about the input errors and output errors, see *Fast Ethernet and Gigabit Ethernet Counters*.

Sample Output

show interfaces statistics (Fast Ethernet)

```
user@host> show interfaces fe-1/3/1 statistics
```

```
Physical interface: fe-1/3/1, Enabled, Physical link is Up
Interface index: 144, SNMP ifIndex: 1042
Description: ford fe-1/3/1
Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags      : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues       : 4 supported, 4 maximum usable queues
Current address: 00:00:5E:00:53:dc, Hardware address: 00:00:5E:00:53:dc
Last flapped    : 2006-04-18 03:08:59 PDT (00:01:24 ago)
Statistics last cleared: Never
```

```

Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
Input errors: 0, Output errors: 0
Active alarms   : None
Active defects  : None
Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
  Flags: SNMP-Traps Encapsulation: ENET2
  Protocol inet, MTU: 1500
    Flags: Is-Primary, DCU, SCU-in

      Packets
Destination class      (packet-per-second)      Bytes
                                (bits-per-second)
      silver1           0                      0
                                (0) (0)
      silver2           0                      0
                                (0) (0)
      silver3           0                      0
                                (0) (0)
Addresses, Flags: Is-Default Is-Preferred Is-Primary
  Destination: 10.27.245/24, Local: 10.27.245.2,
  Broadcast: 10.27.245.255
  Protocol iso, MTU: 1497
  Flags: Is-Primary

```

show interfaces statistics (Gigabit Ethernet PIC—Egress)

user@host> **show interfaces ge-5/2/0 statistics detail**

```

Physical interface: ge-5/2/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 519, Generation: 149
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags        : None
  CoS queues        : 8 supported, 8 maximum usable queues
  Hold-times        : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:74, Hardware address: 00:00:5E:00:53:74
  Last flapped      : 2009-11-11 11:24:00 PST (09:23:08 ago)
  Statistics last cleared: 2009-11-11 17:50:58 PST (02:56:10 ago)
  Traffic statistics:
    Input bytes      :          271524          0 bps
    Output bytes     :        37769598        352 bps

```

```

Input  packets:          3664          0 pps
Output packets:          885790        0 pps
IPv6 transit statistics:
  Input  bytes   :              0
  Output bytes   :         16681118
  Input  packets:              0
  Output packets:         362633
Multicast statistics:
  IPV4 multicast statistics:
    Input  bytes   :         112048          0 bps
    Output bytes   :        20779920          0 bps
    Input  packets:         1801          0 pps
    Output packets:        519498          0 pps
  IPV6 multicast statistics:
    Input  bytes   :         156500          0 bps
    Output bytes   :        16681118          0 bps
    Input  packets:         1818          0 pps
    Output packets:        362633          0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
  L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, 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          882558          882558              0
  1 expedited-fo              0              0              0
  2 assured-forw              0              0              0
  3 network-cont          3232          3232              0
Active alarms   : None
Active defects  : None

Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
  Input  bytes   :         271524
  Output bytes   :        37769598
  Input  packets:          3664
  Output packets:        885790

```

```

IPv6 transit statistics:
  Input  bytes   :                0
  Output bytes   :            16681118
  Input  packets:                0
  Output packets:            362633
Local statistics:
  Input  bytes   :            271524
  Output bytes   :            308560
  Input  packets:            3664
  Output packets:            3659
Transit statistics:
  Input  bytes   :                0                0 bps
  Output bytes   :            37461038            0 bps
  Input  packets:                0                0 pps
  Output packets:            882131                0 pps
IPv6 transit statistics:
  Input  bytes   :                0
  Output bytes   :            16681118
  Input  packets:                0
  Output packets:            362633
Multicast statistics:
  IPV4 multicast statistics:
    Input  bytes   :            112048                0 bps
    Output bytes   :            20779920            0 bps
    Input  packets:            1801                0 pps
    Output packets:            519498                0 pps
  IPV6 multicast statistics:
    Input  bytes   :            156500                0 bps
    Output bytes   :            16681118            0 bps
    Input  packets:            1818                0 pps
    Output packets:            362633                0 pps
Protocol inet, MTU: 1500, Generation: 151, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.40.40.0/30, Local: 10.40.40.2, Broadcast: 10.40.40.3,
Generation: 167
Protocol inet6, MTU: 1500, Generation: 152, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: ::10.40.40.0/126, Local: ::10.40.40.2
Generation: 169
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::21d:b5ff:fe61:d974
Protocol multiservice, MTU: Unlimited, Generation: 171
Generation: 153, Route table: 0

```

```

Policer: Input: __default_arp_policer__

```

show interfaces statistics detail (Aggregated Ethernet)

```

user@host> show interfaces ae0 detail

```

```

Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 186, SNMP ifIndex: 111, Generation: 187
  Link-level type: Ethernet, MTU: 1514, Speed: 2000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5E:0053:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped      : Never
  Statistics last cleared: 2006-12-23 03:04:16 PST (01:16:24 ago)
  Traffic statistics:
    Input  bytes   :                28544                0 bps
    Output bytes   :                39770                0 bps
    Input  packets :                 508                0 pps
    Output packets :                 509                0 pps
    Input  bytes   :             IPv6 28544
    Output bytes   :             IPv6 0
    Input  packets :             IPv6 508
    Output packets :             IPv6 0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

Logical interface ae0.0 (Index 67) (SNMP ifIndex 139) (Generation 145)
  Flags: SNMP-Traps Encapsulation: ENET2
  Statistics          Packets          pps          Bytes          bps
  Bundle:
    Input  :             508             0          28544             0
    Output :             509             0          35698             0
  Link:
    ge-3/3/8.0
      Input  :             508             0          28544             0
      Output :              0              0              0              0
    ge-3/3/9.0

```



```

      Input :           0           0           0           0
      Output:           0           0           0           0
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
ge-3/3/8.0         0             0             0             0
ge-3/3/9.0         0             0             0             0
Egress queues: 8 supported, 8 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
Protocol inet, MTU: 1500, Generation: 166, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
  Generation: 159
Protocol inet6, MTU: 1500, Generation: 163, Route table: 0
Flags: Is-Primary
Addresses, Flags: Is-Preferred
  Destination: fe80::/64, Local: fe80::206:5bff:fe05:c321,
  Broadcast: Unspecified, Generation: 161

```

show interfaces statistics detail (Aggregated Ethernet—Ingress)

user@host> **show interfaces statistics detail ae0 | no-more**

```

Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 504, Generation: 278
  Link-level type: Ethernet, MTU: 1514, Speed: 1Gbps, BPDU Error: None, MAC-REWRITE
  Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5E:00:53:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped   : 2009-11-09 03:30:23 PST (00:01:28 ago)
  Statistics last cleared: 2009-11-09 03:26:18 PST (00:05:33 ago)
Traffic statistics:
  Input bytes   :           544009602           54761856 bps
  Output bytes  :             3396             0 bps
  Input packets:           11826292           148809 pps
  Output packets:             42             0 pps
IPv6 transit statistics:
  Input bytes   :           350818604

```

```

Output bytes : 0
Input packets: 7626488
Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
Ingress 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
Egress queues: 8 supported, 4 in use
Queue counters:
  Queued packets  Transmitted packets  Dropped packets
  0 best-effort    21                 21                0
  1 expedited-fo   0                  0                 0
  2 assured-forw   0                  0                 0
  3 network-cont   451                451               0

Logical interface ae0.0 (Index 70) (SNMP ifIndex 574) (Generation 177)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Statistics
  Packets      pps      Bytes      bps
Bundle:
  Input :      11826292    148809    544009602    54761856
  Output:       42         0         3396         0
Link:
  ge-5/2/0.0
  Input :      11826292    148809    544009602    54761856
  Output:       42         0         3396         0
Marker Statistics:
  Marker Rx    Resp Tx    Unknown Rx    Illegal Rx
  ge-5/2/0.0      0         0             0             0
Protocol inet, MTU: 1500, Generation: 236, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.30.30.0/30, Local: 10.30.30.2, Broadcast: 10.30.30.3,
Generation: 310
Protocol inet6, MTU: 1500, Generation: 237, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: ::10.30.30.0/126, Local: ::10.30.30.2
Generation: 312
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::21d:b5ff:fe61:dbf0

```

```

Protocol multiservice, MTU: Unlimited, Generation: 314
Generation: 238, Route table: 0
Policer: Input: __default_arp_policer__

```

show interfaces statistics detail (Aggregated Ethernet—Egress)

user@host> show interfaces statistics detail ae0 | no-more

```

Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 501, Generation: 319
  Link-level type: Ethernet, MTU: 1514, Speed: 1Gbps, BPDU Error: None, MAC-REWRITE
  Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5E:00:53:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped   : 2009-11-09 03:30:24 PST (00:02:42 ago)
  Statistics last cleared: 2009-11-09 03:26:42 PST (00:06:24 ago)
  Traffic statistics:
    Input  bytes   :                440                0 bps
    Output bytes   :          1047338120          54635848 bps
    Input  packets :                7                0 pps
    Output packets :          22768200          148466 pps
  IPv6 transit statistics:
    Input  bytes   :                288
    Output bytes   :          723202616
    Input  packets :                4
    Output packets :          15721796
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
  0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
  0
  Ingress 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
  Egress queues: 8 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort   201985796          201985796          0

```

```

1 expedited-fo          0          0          0
2 assured-forw          0          0          0
3 network-cont         65         65          0

Logical interface ae0.0 (Index 72) (SNMP ifIndex 505) (Generation 204)
  Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
  Statistics          Packets          pps          Bytes          bps
  Bundle:
    Input :              7              0              440              0
    Output:      22768200      148466      1047338120      54635848
  Link:
    ge-2/1/6.0
      Input :              7              0              440              0
      Output:      22768200      148466      1047338120      54635848
  Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
    ge-2/1/6.0              0              0              0              0
  Protocol inet, MTU: 1500, Generation: 291, Route table: 0
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.30.30.0/30, Local: 10.30.30.1, Broadcast: 10.30.30.3,
  Generation: 420
    Protocol inet6, MTU: 1500, Generation: 292, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::/26, Local: ::10.30.30.1
    Generation: 422
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21f:12ff:fec2:37f0
    Protocol multiservice, MTU: Unlimited, Generation: 424
    Generation: 293, Route table: 0
      Policer: Input: __default_arp_policer__

```

show interfaces statistics (SONET/SDH)

user@host> **show interfaces statistics detail so-3/0/0 | no-more**

```

Physical interface: so-3/0/0, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 538, Generation: 283
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC192,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:

```

```

Input : 13 (last seen 00:00:04 ago)
Output: 14 (last sent 00:00:02 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Opened, iso: Not-configured, mpls: Not-configured

CHAP state: Closed
PAP state: Closed
CoS queues      : 8 supported, 8 maximum usable queues
Last flapped    : 2009-11-09 02:52:34 PST (01:12:39 ago)
Statistics last cleared: 2009-11-09 03:58:54 PST (00:06:19 ago)
Traffic statistics:
Input bytes  :          2559160294          54761720 bps
Output bytes :           10640          48 bps
Input packets:          55633975          148809 pps
Output packets:           216           0 pps
IPv6 transit statistics:
Input bytes  :          647922328
Output bytes :              0
Input packets:          14085269
Output packets:              0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Bucket drops: 0,
Policed discards: 0, L3 incompletes: 0,
L2 channel errors: 0, L2 mismatch timeouts: 0, HS link CRC errors: 0, HS link
FIFO overflows: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, HS link FIFO
underflows: 0, MTU errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
0 best-effort          4              4              0
1 expedited-fo         0              0              0
2 assured-forw         0              0              0
3 network-cont        213             213             0
SONET alarms   : None
SONET defects  : None

Logical interface so-3/0/0.0 (Index 72) (SNMP ifIndex 578) (Generation 182)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
Protocol inet, MTU: 4470, Generation: 244, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.30.30.0/30, Local: 10.30.30.2, Broadcast: 10.30.30.3,
Generation: 322
Protocol inet6, MTU: 4470, Generation: 245, Route table: 0

```

```

Addresses, Flags: Is-Preferred Is-Primary
  Destination: ::10.30.30.0/126, Local: ::10.30.30.2
Generation: 324
Addresses, Flags: Is-Preferred
  Destination: fe80::/64, Local: fe80::2a0:a5ff:fe61:9264
Generation: 326

```

show interfaces statistics (Aggregated SONET/SDH—Ingress)

user@host> **show interfaces statistics detail as0 | no-more**

```

Physical interface: as0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 534, Generation: 282
  Link-level type: PPP, MTU: 4474, Speed: OC192, Minimum links needed: 1, Minimum
bandwidth needed: 0
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags        : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Last flapped      : 2009-11-09 03:45:53 PST (00:09:38 ago)
  Statistics last cleared: 2009-11-09 03:48:17 PST (00:07:14 ago)
  Traffic statistics:
    Input  bytes   :           2969786332           54761688 bps
    Output bytes   :             11601             0 bps
    Input  packets:           64560636           148808 pps
    Output packets:             225             0 pps
  IPv6 transit statistics:
    Input  bytes   :           2086013152
    Output bytes   :             0
    Input  packets:           45348114
    Output packets:             0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
  Egress queues: 8 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort    3              3              0
    1 expedited-fo   0              0              0
    2 assured-forw    0              0              0
    3 network-cont   222            222            0

```

```

Logical interface as0.0 (Index 71) (SNMP ifIndex 576) (Generation 179)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
  Statistics          Packets          pps          Bytes          bps
Bundle:
  Input :             64560550         148808         2969785300         54761688
  Output:              139              0             10344              0
Link:
  so-3/0/0.0
  Input :             64560550         148808         2969785300         54761688
  Output:              139              0             10344              0
Protocol inet, MTU: 4470, Generation: 240, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.30.30.0/30, Local: 10.30.30.2, Broadcast: 10.30.30.3,
Generation: 316
Protocol inet6, MTU: 4470, Generation: 241, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: ::10.30.30.0/126, Local: ::10.30.30.2
Generation: 318
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::2a0:a5ff:fe61:9264
Generation: 320

```

show interfaces statistics (Aggregated SONET/SDH—Egress)

user@host> **show interfaces statistics detail as0 | no-more**

```

Physical interface: as0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 565, Generation: 323
  Link-level type: PPP, MTU: 4474, Speed: OC192, Minimum links needed: 1, Minimum
bandwidth needed: 0
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags       : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Last flapped     : 2009-11-09 03:43:37 PST (00:12:48 ago)
  Statistics last cleared: 2009-11-09 03:48:54 PST (00:07:31 ago)
  Traffic statistics:
    Input bytes   :             11198             392 bps
    Output bytes  :          3101452132          54783448 bps
    Input packets:              234              0 pps
    Output packets:         67422937         148868 pps
  IPv6 transit statistics:
    Input bytes   :             5780
    Output bytes  :          2171015678

```

```

    Input  packets:                72
    Output packets:             47195993
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort      67422830          67422830              0
  1 expedited-fo           0                0                  0
  2 assured-forw           0                0                  0
  3 network-cont         90               90                  0

Logical interface as0.0 (Index 71) (SNMP ifIndex 548) (Generation 206)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :          144          0        10118        392
  Output:      67422847      148868    3101450962    54783448
Link:
  so-0/1/0.0
    Input :          144          0        10118        392
    Output:      67422847      148868    3101450962    54783448
Protocol inet, MTU: 4470, Generation: 295, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.30.30.0/30, Local: 10.30.30.1, Broadcast: 10.30.30.3,
Generation: 426
Protocol inet6, MTU: 4470, Generation: 296, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: ::/26, Local: ::10.30.30.1
Generation: 428
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::2a0:a5ff:fe63:1d0a
Generation: 429

```

show interfaces statistics (MX Series Routers)

user@host> show interfaces xe-0/0/0 statistics

```

Physical interface: xe-0/0/0, Enabled, Physical link is Up
Interface index: 145, SNMP ifIndex: 592
Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:

```



```

None, Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Enabled
Device flags      : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags        : None
CoS queues        : 8 supported, 8 maximum usable queues
Current address: 00:00:5E:00:53:f0, Hardware address: 00:00:5E:00:53:f0
Last flapped      : 2013-10-26 03:20:40 test (2w3d 03:29 ago)
Statistics last cleared: Never
Input rate        : 0 bps (0 pps)
Output rate       : 0 bps (0 pps)
Input errors: 0, Output errors: 0
Active alarms     : LINK
Active defects    : LINK
PCS statistics
  Bit errors              Seconds
  109
  Errored blocks          109
Interface transmit statistics: Disabled

```

show interfaces statistics (MX Series Routers: Dynamic Interfaces with RPF Check Detail)

user@host> **show interfaces statistics pp0.3221225475 detail**

```

Logical interface pp0.3221225475(Index 536870921)(SNMP ifIndex 200000009)
(Generation 6)
  Flags: Up Point-To-Point Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 1,
    Session AC name: B, Remote MAC address:00:00:5E:00:53:01,
    Underlying interface: xe-1/0/0.3221225474 (Index 536870919)
    Ignore End-Of-List tag: Disable
  Bandwidth: 0
  Traffic statistics:
    Input  bytes   :                34
    Output bytes   :                 0
    Input  packets :                 1
    Output packets :                 1
  Local statistics:
    Input  bytes   :                 0
    Output bytes   :                 0
    Input  packets :                 0
    Output packets :                 0
  Transit statistics:
    Input  bytes   :                34
                                0 bps

```

```

        Output bytes      :                0                0 bps
        Input  packets:      1                0 pps
        Output packets:      1                0 pps
    Keepalive settings: Interval 30 seconds, Up-count 3, Down-count 3
    LCP state: Opened
    NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
    CHAP state: Success
    PAP state: Closed
    Protocol inet, MTU: 1492
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
NH drop cnt: 0
    Generation: 0, Route table: 0
    Flags: uRPF, Unnumbered
    RPF Failures: Packets: 0, Bytes: 0
    Donor interface: lo0.0 (Index 320)
    Input Filters: upstrml-inet-pp0.3221225475-in
    Output Filters: dwnstrml-inet-pp0.3221225475-out
    Addresses, Flags: Is-Primary
        Destination: Unspecified, Local: 10.255.96.19, Broadcast: Unspecified,
Generation: 0

```

show interfaces statistics (PTX Series Packet Transport Routers)

user@host> show interfaces statistics em0

```

Physical interface: em0, Enabled, Physical link is Up
    Interface index: 8, SNMP ifIndex: 0
    Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
    Device flags      : Present Running
    Interface flags: SNMP-Traps
    Link type         : Full-Duplex
    Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
    Last flapped      : Never
    Statistics last cleared: Never
Input packets : 212620
Output packets: 71
    Input errors: 0, Output errors: 0

    Logical interface em0.0 (Index 3) (SNMP ifIndex 0)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 212590
    Output packets: 71
    Protocol inet, MTU: 1500

```

```

Flags: Is-Primary
Addresses, Flags: Is-Default Is-Preferred Is-Primary
  Destination: 192.168.3/24, Local: 192.168.3.30,
  Broadcast: 192.168.3.255

```

show interfaces statistics (ACX Series routers)

user@host> **show interfaces statistics ge-0/1/7**

```

Physical interface: ge-0/1/7, Enabled, Physical link is Down
  Interface index: 151, SNMP ifIndex: 524
  Link-level type: Ethernet, Media type: Copper, MTU: 1514, Link-mode: Full-duplex,
  Speed: 1000mbps, BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags      : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags       : None
  CoS queues       : 8 supported, 8 maximum usable queues
  Current address: 00:00:5E:00:53:a3, Hardware address: 00:00:5E:00:53:a3
  Last flapped    : 2012-05-11 04:25:28 PDT (2d 20:23 ago)
  Statistics last cleared: 2012-05-13 23:07:23 PDT (01:41:25 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms   : LINK
  Active defects  : LINK
  Interface transmit statistics: Disabled

```

show interfaces targeting (Aggregated Ethernet for Subscriber Management)

Syntax

```
show interfaces targeting aex
```

Release Information

Command introduced in Junos OS Release 11.2.

Description

(MX Series routers only) Display status information about the distribution of subscribers on different links in an aggregated Ethernet bundle.

Required Privilege Level

view

Output Fields

[Table 70 on page 1422](#) lists the output fields for the **show interfaces targeting** command. Output fields are listed in the approximate order in which they appear.

Table 70: show interfaces targeting Output Fields

Field Name	Field Description	Level of Output
Aggregated Ethernet Interface		
Aggregated interface	Name of the aggregated Ethernet bundle.	All levels
Redundancy mode	Redundancy mechanism on the interface: Link Level Redundancy or FPC Redundancy .	All levels
Total number of distributed interfaces	Number of distributed links in the bundle.	All levels
Physical Interface		
Physical interface	Name of the physical interface and state of the interface.	All levels
Link status	Status of the link on the physical interface: up or down .	
Number of primary distributions	Number of subscribers distributed on primary links.	All levels

Table 70: show interfaces targeting Output Fields (*continued*)

Field Name	Field Description	Level of Output
Number of backup distributions	Number of subscribers distributed on backup links.	All levels

Sample Output

show interfaces targeting ae0

user@host> **show interfaces targeting ae0**

```

Aggregated interface: ae0
Redundancy mode: Link Level Redundancy
Total number of distributed interfaces: 3
Physical interface: ge-1/0/0, Link status: Up
Number of primary distributions: 200
Number of backup distributions: 200
Physical interface: ge-1/1/0, Link status: Up
Number of primary distributions: 200
Number of backup distributions: 199
Physical interface: ge-2/0/7, Link status: Up
Number of primary distributions: 200
Number of backup distributions: 200
Physical interface: ge-2/0/8, Link status: Up
Number of primary distributions: 199
Number of backup distributions: 200

```

show mld group

Syntax

```
show mld group
<brief | detail>
<group-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display information about Multicast Listener Discovery (MLD) group membership.

Options

none—Display standard information about all MLD groups.

brief | detail—(Optional) Display the specified level of output.

group-name—(Optional) Display MLD information about the specified group.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear mld membership](#) | 1265

List of Sample Output

[show mld group \(Include Mode\) on page 1425](#)

[show mld group \(Exclude Mode\) on page 1426](#)

[show mld group brief on page 1427](#)

[show mld group detail \(Include Mode\) on page 1427](#)

[show mld group detail \(Exclude Mode\) on page 1428](#)

Output Fields

[Table 71 on page 1425](#) describes the output fields for the **show mld group** command. Output fields are listed in the approximate order in which they appear.

Table 71: show mld group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the MLD membership report; local means that the local router joined the group itself.	All levels
Group	Group address.	All levels
Source	Source address.	All levels
Group Mode	Mode the SSM group is operating in: Include or Exclude .	All levels
Last reported by	Address of the host that last reported membership in this group.	All levels
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	detail
Timeout	Time remaining until the group membership is removed.	brief none
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	detail
Type	Type of group membership: <ul style="list-style-type: none"> • Dynamic—Host reported the membership. • Static—Membership is configured. 	All levels

Sample Output

show mld group
(Include Mode)

user@host> **show mld group**

```

Interface: fe-0/1/2.0
  Group: ff02::1:ff05:1a67
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      245 Type: Dynamic
  Group: ff02::1:ffa8:c35e

```

```

    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      241 Type: Dynamic
Group: ff02::2:43e:d7f6
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      244 Type: Dynamic
Group: ff05::2
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      244 Type: Dynamic
Interface: local
    Group: ff02::2
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic
Group: ff02::16
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic

```

show mld group (Exclude Mode)

user@host> show mld group

```

Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
    Group: ff02::6
    Source: ::
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Timeout:      245 Type: Dynamic
Group: ff02::16
    Source: ::
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Timeout:      28 Type: Dynamic
Interface: local
    Group: ff02::2
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic
Group: ff02::16

```



```

Source: ::
Last reported by: Local
Timeout:      0 Type: Dynamic

```

show mld group brief

The output for the **show mld group brief** command is identical to that for the **show mld group** command. For sample output, see [show mld group \(Include Mode\) on page 1425](#) [show mld group \(Exclude Mode\) on page 1426](#).

show mld group detail (Include Mode)

```
user@host> show mld group detail
```

```

Interface: fe-0/1/2.0
  Group: ff02::1:ff05:1a67
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      224 Type: Dynamic
  Group: ff02::1:ffa8:c35e
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      220 Type: Dynamic
  Group: ff02::2:43e:d7f6
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      223 Type: Dynamic
  Group: ff05::2
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      223 Type: Dynamic
Interface: so-1/0/1.0
  Group: ff02::2
    Group mode: Include
    Source: ::
    Last reported by: fe80::280:42ff:fe15:f445
    Timeout:      258 Type: Dynamic
Interface: local
  Group: ff02::2
    Group mode: Include

```

```

    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic
Group: ff02::16
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic

```

show mld group detail (Exclude Mode)

user@host> **show mld group detail**

```

Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
  Group: ff02::6
    Group mode: Exclude
    Source: ::
    Source timeout: 0
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Group timeout:    226 Type: Dynamic
  Group: ff02::16
    Group mode: Exclude
    Source: ::
    Source timeout: 0
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Group timeout:    246 Type: Dynamic
Interface: local
  Group: ff02::2
    Group mode: Exclude
    Source: ::
    Source timeout: 0
    Last reported by: Local
    Group timeout:    0 Type: Dynamic
  Group: ff02::16
    Group mode: Exclude
    Source: ::
    Source timeout: 0
    Last reported by: Local
    Group timeout:    0 Type: Dynamic

```

show mld interface

Syntax

```
show mld interface  
<brief | detail>  
<interface-name>  
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display information about multipoint Listener Discovery (MLD)-enabled interfaces.

Options

none—Display standard information about all MLD-enabled interfaces.

brief | detail—(Optional) Display the specified level of output.

interface-name—(Optional) Display information about the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear mld membership](#) | 1265

List of Sample Output

[show mld interface on page 1432](#)

[show mld interface brief on page 1432](#)

[show mld interface detail on page 1432](#)

[show mld interface <interface-name> on page 1433](#)

Output Fields

[Table 72 on page 1430](#) describes the output fields for the **show mld interface** command. Output fields are listed in the approximate order in which they appear.

Table 72: show mld interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the router that has been elected to send membership queries.	All levels
State	State of the interface: Up or Down .	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the interface.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy at the MLD interface.	All levels
Timeout	How long until the MLD querier is declared to be unreachable, in seconds.	All levels
Version	MLD version being used on the interface: 1 or 2 .	All levels
Groups	Number of groups on the interface.	All levels
Passive	<p>State of the passive mode option:</p> <ul style="list-style-type: none"> • On—Indicates that the router can run IGMP or MLD on the interface but not send or receive control traffic such as IGMP or MLD reports, queries, and leaves. • Off—Indicates that the router can run IGMP or MLD on the interface and send or receive control traffic such as IGMP or MLD reports, queries, and leaves. <p>The passive statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the on state declaration:</p> <ul style="list-style-type: none"> • send-general-query—The interface sends general queries. • send-group-query—The interface sends group-specific and group-source-specific queries. • allow-receive—The interface receives control traffic 	All levels
OIF map	Name of the OIF map associated to the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map used on the interface, if configured.	All levels
Group limit	Maximum number of groups allowed on the interface. Any memberships requested after the limit is reached are rejected.	All levels

Table 72: show mld interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Group threshold	Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.	All levels
Group log-interval	Time (in seconds) between consecutive log messages.	All levels
Immediate Leave	State of the immediate leave option: <ul style="list-style-type: none"> • On—Indicates that the router removes a host from the multicast group as soon as the router receives a multicast listener done message from a host associated with the interface. • Off—Indicates that after receiving a multicast listener done message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds. 	All levels
Distributed	State of MLD, which, by default, takes place on the Routing Engine for MX Series routers but can be distributed to the Packet Forwarding Engine to provide faster processing of join and leave events. <ul style="list-style-type: none"> • On—distributed MLD is enabled. 	All levels
Configured Parameters	Information configured by the user. <ul style="list-style-type: none"> • MLD Query Interval (.1 secs)—Interval at which this router sends membership queries when it is the querier. • MLD Query Response Interval (.1 secs)—Time that the router waits for a report in response to a general query. • MLD Last Member Query Interval (.1 secs)—Time that the router waits for a report in response to a group-specific query. • MLD Robustness Count—Number of times the router retries a query. 	All levels
Derived Parameters	Derived information. <ul style="list-style-type: none"> • MLD Membership Timeout (.1 secs)—Timeout period for group membership. If no report is received for these groups before the timeout expires, the group membership will be removed. • MLD Other Querier Present Timeout (.1 secs)—Time that the router waits for the IGMP querier to send a query. 	All levels

Sample Output

show mld interface

user@host> show mld interface

```

Interface: fe-0/0/0
  Querier: None
  State: Up          Timeout:      0    Version:  1    Groups:      0
  SSM Map Policy: ssm-policy-A
Interface: at-0/3/1.0
  Querier: 8038::c0a8:c345
  State: Up          Timeout:    None    Version:  1    Groups:      0
  SSM Map Policy: ssm-policy-B
Interface: fe-1/0/1.0
  Querier: ::192.168.195.73
  State: Up          Timeout:    None    Version:  1    Groups:      3
  SSM Map Policy: ssm-policy-C
  SSM map: ipv6map1
Immediate Leave: On

Promiscuous Mode: Off
Passive: Off
Distributed: OnConfigured Parameters:

Configured Parameters:
MLD Query Interval (.1 secs): 1250
MLD Query Response Interval (.1 secs): 100
MLD Last Member Query Interval (.1 secs): 10
MLD Robustness Count: 2

Derived Parameters:
MLD Membership Timeout (.1secs): 2600
MLD Other Querier Present Timeout (.1 secs): 2550

```

show mld interface brief

The output for the **show mld interface brief** command is identical to that for the **show mld interface** command. For sample output, see [show mld interface on page 1432](#).

show mld interface detail

The output for the **show mld interface detail** command is identical to that for the **show mld interface** command. For sample output, see [show mld interface on page 1432](#).

show mld interface <interface-name>

user@host# **show mld interface ge-3/2/0.0**

```
Interface: ge-3/2/0.0
Querier: 203.0.113.111
State: Up Timeout:    None Version:  3 Groups:      1
Group limit: 8
Group threshold: 60
Group log-interval: 10
Immediate leave: Off
Promiscuous mode: Off   Distributed: On
```

show network-access aaa subscribers session-id

Syntax

```
show network-access aaa subscribers session-id session-id
<brief | detail>
```

Release Information

Command introduced in Junos OS Release 10.0.

Description

Display information about the specified subscriber session.

Options

session-id—ID of the subscriber session.

brief | detail—(Optional) Display the specified level of information.

Required Privilege Level

view

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

[Local and Remote Service Activation and Deactivation Using the CLI | 21](#)

[Deactivating a Single Instance of a Subscriber Service | 28](#)

[Deactivating All Instances of a Subscriber Service | 31](#)

[Verifying Subscriber Services with Multiple Instances | 34](#)

Policy and Charging Enforcement Function Overview for Broadband Wireline Subscribers

Understanding Gy Interactions Between the Router and the OCS

Understanding Interactions Between the PCRF, PCEF, and OCS

List of Sample Output

[show network-access aaa subscribers session-id brief on page 1438](#)

[show network-access aaa subscribers session-id 2 \(Flat File Accounting\) on page 1439](#)

[show network-access aaa subscribers session-id detail on page 1439](#)

[show network-access aaa subscribers session-id detail \(Service with Multiple Instances\) on page 1440](#)

[show network-access aaa subscribers session-id detail \(Single Session Dual Stack with active V4 and V6 subscribers\) on page 1441](#)

Output Fields

Table 73 on page 1435 lists the output fields for the **show network-access aaa subscribers session-id** command. Output fields are listed in the approximate order in which they appear.

Table 73 on page 1435 lists the output fields for the **show network-access aaa subscribers session-id** command. Output fields are listed in the approximate order in which they appear.

Table 73: show network-access aaa subscribers session-id Output Fields

Field Name	Field Description	Level of Output
Type and Client type	Type of client.	All levels
Accounting	Status of the accounting configuration for the service, on or off , and the type of accounting, time , volume+time , or flat-file . The time and volume+time types are configured in RADIUS Service-Statistics VSA [26-69].	brief none
Service type	Type of accounting: volume , time , volume+time , or na .	brief
Quota	Quota for service: volume (in Mbps) or time (seconds).	brief
Username	Name of the user logged in to the session.	detail
Stripped username	Username after the domain has been removed.	detail
Logical system/Routing instance and AAA Logical system/Routing instance	Name of the routing instance, logical system name, or both used for the session.	All levels
Target Logical system/Routing instance	Logical system/routing instance to which the session is mapped.	detail
Access-profile	Access profile used for AAA services for the session.	detail
Session ID	ID of the subscriber session. The session ID value displayed under Service name is the service session ID.	detail

Table 73: show network-access aaa subscribers session-id Output Fields (*continued*)

Field Name	Field Description	Level of Output
Accounting Session ID	ID of the accounting session (RADIUS attribute 44). The ID appears in decimal or description format, as specified by the accounting-session-id-format statement.	detail
Multi Accounting Session ID	Bundle ID for MLPPP sessions. Acct-Multi-Session-Id (RADIUS attribute 50) uses the value of the session database bundle session ID to enable RADIUS to link together multiple related sessions. The value of this field is zero when no MLPPP sessions exist.	detail
IP Address	IP address of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
IPv6 Address	IPv6 address of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
IPv6 Prefix	IPv6 prefix of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
Authentication State	State of the subscriber authentication session: AuthInit , AuthStart , AuthChallenge , AuthRedirect , AuthClntRespWait , AuthAcctVolStatsAckWait , AuthAcctStopAckWait , AuthServCreateRespWait , AuthLogoutStart , AuthStateActive , AuthClntLogoutRespWait , AuthProfileUpdateWait , AuthProvisionRespWait , AuthProvisionServiceCreationWait	detail
Gx-Plus Provisioning State	State of Gx-Plus provisioning: <ul style="list-style-type: none"> • ignored—Subscriber has no IPv4 address or NAS-Port-ID. • in-progress—Provisioning is in progress. • logout—Subscriber logout is in progress. • logout-done—Logout response has been received. • response-received—Provisioning response has been received. 	detail
Pcrf Provisioning State	State of PCRF provisioning: <ul style="list-style-type: none"> • active—PCRF provisioning is active. • ignored—Subscriber has no IPv4 address or NAS-Port-ID. • in-progress—Provisioning is in progress. • logout—Subscriber logout is in progress. • logout-done—Logout response has been received. • response-received—Provisioning response has been received. 	detail

Table 73: show network-access aaa subscribers session-id Output Fields (*continued*)

Field Name	Field Description	Level of Output
Pcrf Subscription-Id-Type	Type of subscriber for a PCRF partition. You can define your own or use a predefined value: 0 (END_USER_E164) , 1 (END_USER_IMSI) , 2 (END_USER_SIP_URI) , 3 (END_USER_NAI) , 4 (END_USER_PRIVATE) .	detail
Pcrf Subscription-Id-Data	Subscriber data string concatenated from a list of user-selected data options used to identify the subscriber type for a PCRF partition; for example, demux0	detail
Ocs Subscription-Id-Type	Type of subscriber for an OCS partition. You can define your own or use a predefined value: 0 (END_USER_E164) , 1 (END_USER_IMSI) , 2 (END_USER_SIP_URI) , 3 (END_USER_NAI) , 4 (END_USER_PRIVATE) .	detail
Ocs Subscription-Id-Data	Subscriber data string concatenated from a list of user-selected data options used to identify the subscriber type for an OCS partition; for example: test-sid	detail
Ocs Interrogation State	State of the OCS interrogation: first , intermediate , final .	detail
Ocs Data State	State of the OCS data: none	detail
Accounting State	State of the subscriber accounting session: Acc-Init , Acc-Start-Sent , Imm-Update-Stats-Pending , Acc-Interim-Sent , Acc-Stop-Stats-Pending , Acc-Stop-Sent , Acc-Stop-On-Fail-Deny-Sent , Acc-Stop-Ackd	detail
Provisioning-type	Provisioning type for this session: <ul style="list-style-type: none"> • gx-plus—Subscriber service uses Gx-Plus provisioning. • jsrc—Subscriber service uses JSRC provisioning. • none—Provisioning is not enabled. 	detail
Service name	Name of the attached service or policy. <ul style="list-style-type: none"> • For RADIUS-activated and CLI-activated services, displays the full activation string for the service. If the activation string includes service parameters, then both the service name and service parameters are displayed. • For JSRC-activated policies—displays the policy name. 	All levels
Service State	State of the service provided in the subscriber session.	detail
Service Family	Network family of the service provided in the subscriber session.	detail

Table 73: show network-access aaa subscribers session-id Output Fields (continued)

Field Name	Field Description	Level of Output
Service Activation Source	Source used to activate the service.	detail
Session uptime	How long the session has been up, in <i>HH:MM:SS</i> .	All levels
Service CC-Service-Identifier	Data identification element of the 3GPP Diameter credit control service charging system that uniquely defines the CC-Service-Context .	detail
Service Rating-Group	Value associated with a charging rule and part of the accounting data stream for the PCRF.	detail
Ocs Control	Whether OCS controls the service: yes or no .	detail
Accounting status	Status of the accounting configuration for the service, on or off , and the type of accounting, time , volume+time , or flat-file . The time and volume+time types are configured in RADIUS Service-Statistics VSA [26-69].	detail
Service accounting session ID	ID of the service accounting session; RADIUS Acct-Session-Id attribute (44). The ID appears in decimal or description format, as specified by the accounting-session-id-format statement.	detail
Service accounting state	State of the service accounting session: Acc-Init , Acc-Start-Sent , Imm-Update-Stats-Pending , Acc-Interim-Sent , Acc-Stop-Stats-Pending , Acc-Stop-Sent , Acc-Stop-On-Fail-Deny-Sent , Acc-Stop-Ackd	detail
Accounting interim interval	Amount of time between interim accounting updates for this service, in seconds; RADIUS Service-Interim-Acct-Interval VSA [26-140] or Diameter Acct-Interim-Interval AVP (85).	detail

Sample Output

show network-access aaa subscribers session-id brief

```
user@host> show network-access aaa subscribers session-id 6 brief
```

Logical system/Routing instance	Client type	Session uptime	Accounting
default:default	dhcp	00:01:29	on/time
Service name	Service type	Quota	Accounting
filter-service	-na-	-na-	off
filter-service-2	volume+time	77.00MB/120secs	off
1337994190863204450	-na-	-na-	off

show network-access aaa subscribers session-id 2 (Flat File Accounting)

```
user@host> show network-access aaa subscribers session-id 2
```

Logical system/Routing instance	Client type	Session-ID	Session uptime
Accounting			
default:default	dhcp	2	00:00:48
on/volume+time			
Service name	Service type	Quota	Accounting
filter-service	-na-	-na-	on/flat-file

show network-access aaa subscribers session-id detail

```
user@host> show network-access aaa subscribers session-id 5 detail
```

```
Type: dhcp
Username: user23@example.net
Stripped username: user23
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:retail-onlinecompany-ca
Access-profile:retailer-onlinecompany-sjc
Session ID: 5
Accounting Session ID: jnpr ge-1/0/0.101:1
Multi Accounting Session ID: 0
IP Address: 192.168.44.104
Authentication State: AuthStateActive
Pcrf session-stamp: 0
Pcrf Provisioning State: active
Pcrf Subscription-Id-Type: 4
Pcrf Subscription-Id-Data: demux0
Ocs Subscription-Id-Type: 15
Ocs Subscription-Id-Data: test-sid
Ocs Interrogation State: intermediate
Ocs Data State: none
```

```

Gx-Plus Provisioning State: response-received
Accounting State: Acc-Interim-Sent
Provisioning-type: jsrsc
Service name: filter-service-1
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: PCRF-LOGIN
  Session ID: 7
  Session uptime: 00:01:33
  Service CC-Service-Identifier: 777
  Service Rating Group: 10
  Ocs Control: yes
Service name: filter-service-2
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: PCRF-LOGIN Session ID: 8
  Session uptime: 00:01:33
  Service CC-Service-Identifier: 778
  Service Rating Group: 11
  Ocs Control: no
Accounting status: on/volume+time
  Service accounting session ID: 1:2-1322506006
  Service accounting state: Acc-Interim-Sent
  Accounting interim interval: 600

```

show network-access aaa subscribers session-id detail (Service with Multiple Instances)

user@host> show network-access aaa subscribers session-id 6 detail

```

Type: dhcp
Stripped username: user-test-fms2
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: attr_test_profile1
Session ID: 6
Accounting Session ID: 6
Multi Accounting Session ID: 0
IP Address: 198.51.100.10
Authentication State: AuthStateActive
Pcrf session-stamp: 0
Accounting State: Acc-Interim-Sent
Provisioning Type: None
Service name: economy-service(up-filter,down-filter)
  Service State: SvcActive

```

```

Service Family: inet
Service Activation Source: Radius at Reauth
Session ID: 7
Session uptime: 00:04:36
Accounting status: on/volume+time
Service accounting session ID: 6:7-1354811427
Service accounting state: Acc-Start-Sent
Accounting interim interval: 600
Service name: economy-service(upstrm-filter,dwnstrm-filter)
Service State: SvcActive
Service Family: inet
Service Activation Source: Radius
Session ID: 8
Session uptime: 00:04:36
Accounting status: on/volume+time
Service accounting session ID: 6:8-1354811427
Service accounting state: Acc-Start-Sent
Accounting interim interval: 600

```

show network-access aaa subscribers session-id detail (Single Session Dual Stack with active V4 and V6 subscribers)

user@host> **show network-access aaa subscribers session-id 26 detail**

```

Type: dhcp
Stripped username: user-test-25
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.10.0.6
IPv6 Address: 00:00:5E:00:53:02
IPv6 Prefix: 00:00:5E:00:53:00/64
Authentication State: AuthStateActive
Pcrf session-stamp: 0
Accounting State: Acc-Interim-Sent
Provisioning Type: None

```

show services analytics agent

Syntax

```
show services analytics agent
<brief | detail>
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Command introduced in Junos OS Release 18.4R1 on PTX Series routers.

Description

Display information about running instances of Network Telemetry Framework (NTF) agent.

Options

none—(Same as brief) Display summary information about analytics agents.

brief | detail—(Optional) Display information about analytics agents for the specified level of output.

Required Privilege Level

view

RELATED DOCUMENTATION

[IPFIX Mediation on the BNG | 667](#)

Configuring NTF Agent

List of Sample Output

[show services analytics agent on page 1443](#)

[show services analytics agent \(Brief\) on page 1443](#)

[show services analytics agent \(Detail\) on page 1444](#)

Output Fields

[Table 74 on page 1442](#) lists the output fields for the **show services analytics agent** command. Output fields are listed in the approximate order in which they appear.

Table 74: show services analytics agent Output Fields

Field Name	Field Description	Level of Output
Agent ID	Name of the agent.	brief none
Output Plugins	Number of output plug-ins configured for the agent.	brief none

Table 74: show services analytics agent Output Fields (continued)

Field Name	Field Description	Level of Output
Input Plugins	Number of input plug-ins configured for the agent.	brief none
Process ID	Number that uniquely identifies the active process for the service agent at the brief and none levels. At the detail level, the process ID is displayed for the analytics agent (the parent NTF agent) and for the active service agents.	All levels
Analytics agent	Information about the parent NTF agent.	detail
Configuration File	Path where the NTF agent configuration file is located.	detail
Log File	Path where logs are stored for the NTF agent.	detail
Service Agent Count	Number of active service agents.	detail
Analytics Service agent(s)	Information about the active service agents.	detail
Agent Name	Name of the service agent.	detail
Input Plugin/s	Name of all input plug-ins configured for the service agent.	detail
Output Plugin/s	Name of all output plug-ins configured for the service agent.	detail

Sample Output

show services analytics agent

```
user@host> show services analytics agent
```

Agent ID	Output Plugins	Input Plugins	Process ID
ipfix	1	2	8368

show services analytics agent (Brief)

```
user@host> show services analytics agent brief
```

Agent ID	Output Plugins	Input Plugins	Process ID
ipfix	1	2	8368

show services analytics agent (Detail)

user@host> show services analytics agent detail

```
Analytics agent:
Process ID           : 6246
Configuration File   : /var/etc/ntf-agent.conf
Log File             : /var/log/ntf-agent.log
Service Agent Count  : 1
Analytics service agent(s):
  Agent Name         : ipfix
  Input Plugin/s     : input-ipfix
  Output Plugin/s    : output-ipfix
  Process ID        : 8368
```

show remote-device-management service-devices

Syntax

```
show remote-device-management service-devices
  name
  <detail>
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Display information about all remote service devices or a specific remote service device.

Options

none—Display summary information about all remote service devices.

detail—Display detailed information about remote service devices.

name—Name of the remote service device.

Required Privilege Level

view

RELATED DOCUMENTATION

[show remote-device-management services](#) | 1453

[show remote-device-management statistics](#) | 1456

[show remote-device-management subscribers](#) | 1461

[show remote-device-management summary](#) | 1465

List of Sample Output

[show remote-device-management service-devices](#) on page 1450

[show remote-device-management service-devices \(Device Name\)](#) on page 1450

[show remote-device-management service-devices \(Device Name Detail\)](#) on page 1450

[show remote-device-management service-devices \(Device Name Extensive\)](#) on page 1451

Output Fields

[Table 75 on page 1446](#) lists the output fields for the **show remote-device-management service-devices** command. Output fields are listed in the approximate order in which they appear.

Table 75: show remote-device-management service-devices Output Fields

Field Name	Field Description	Level of Output
Device Name	Name of the remote service device.	All levels
State	State —State of the remote device service management (RDSM) connection to the service device, Up or Down .	All levels
Eligible Services	<p>Number of remote subscriber services that are eligible to be provisioned on the device.</p> <p>When a remote device is up without interruption, the Eligible Services and Provisioned Services counters should be equal.</p> <p>If the Provisioned Services counter is less than the Eligible Services counter, you may need to reconfigure the remote device to provision the outstanding eligible services.</p>	All levels
Provisioned Services	Number of remote subscriber services that are provisioned on the device.	All levels
Address	IP address of the remote device used to configure the subscriber service; unique across all routing instances.	detail extensive
Last State Change Time	Timestamp when the RDSM connection to the remote device last changed state.	detail extensive
Vlan Id List	List of VLAN ranges and IDs that are served by the remote device and make up the access domain. The access domain corresponds to the set of subscriber-facing Layer 2 locations that map to the device.	detail extensive
Dictionary File	The absolute file path on the router for the vendor-specific dictionary that defines the set of NETCONF XML protocol commands required to provision, deprovision, and roll back a subscriber service on the remote device.	detail extensive

Table 75: show remote-device-management service-devices Output Fields (continued)

Field Name	Field Description	Level of Output
Provisioning Method	<p>Attributes configured for the NETCONF XML Protocol method for provisioning and deprovisioning services on the remote device:</p> <ul style="list-style-type: none"> • User Name—Name used to access the remote device during service management. • Connection Retry Interval—The interval between successive attempts to establish a NETCONF session with the remote device. • Response Timeout—Period during which the device must respond to an attempt to provision or deprovision a service. • Response Timeout Count—Number of consecutive response timeouts that occur before the BNG takes action. • Bulk Interval—Interval during which multiple services are provisioned or deprovisioned based on the assigned dictionary before the configuration is committed to the service device. • Bulk Limit—Maximum number of services provisioned or deprovisioned based on the assigned dictionary during the bulk interval—before the configuration is committed to the service device • Reconfigure Bulk Limit—Maximum number of services provisioned on the service device for the access domain when the device is reconfigured; number is based on the assigned dictionary before the configuration is committed to the service device • Port—TCP port number for the NETCONF protocol session. 	detail extensive
Round Trip Time (millisec)	<p>Aggregate duration for all dictionary RPCs required to provision or deprovision a service: Minimum, Maximum, Average, and Last.</p> <p>When the bulk limit is set to 1, these statistics are reported per remote service.</p> <p>When the bulk limit is greater than 1, the duration includes all services that are part of the bulk configuration.</p>	detail extensive
Bulk Count	<p>Number of services included in a bulk configuration: Minimum, Maximum, Average, and Last. The number is determined by the bulk-limit and bulk-interval configuration for the remote device at the [edit system services remote-device-management service-device device-name provisioning-method netconf] hierarchy level.</p> <p>When the bulk limit is set to 1, the count is 1 for all fields.</p>	detail extensive
Service Session ID	ID number for the remote service session.	extensive

Table 75: show remote-device-management service-devices Output Fields (*continued*)

Field Name	Field Description	Level of Output
Subscriber Session ID	ID number for the subscriber session associated with the service session.	extensive
State	<p>Status of the service session; represents the RDSM processing state:</p> <ul style="list-style-type: none"> • Active—Remote service session is up and fully provisioned. • Deprovisioning—Remote service session is in the process of being deprovisioned. • Init—Transient state that indicates that the state is transitioning to the provisioning state. • Provisioning—Remote service session is in the process of being provisioned. • Provisioning-Complete—Transient state that indicates service provisioning has completed on the device; transitions to Active or Provisioning-Rollback. • Provisioning-Rollback—Rollback of the service provisioning on the device is in progress. Provisioning is rolled back when provisioning fails for one or more other eligible service devices. 	extensive
Service Name	Name of the remote service.	extensive

Table 75: show remote-device-management service-devices Output Fields (continued)

Field Name	Field Description	Level of Output
Service Provisioning	<p>Status counts for all provisioning actions attempted for the remote device.</p> <ul style="list-style-type: none"> • Attempted—Number of service provisioning attempts for a service device in the Up state. • Succeeded—Number of services successfully provisioned. • Failed—Number of services that failed provisioning for all reasons <ul style="list-style-type: none"> • Reject—Number of failures due to an explicit error response during provisioning. • Timeout—Number of failures due to no response during provisioning. • In Progress—Number of services being provisioned but not yet completed. • Reconfigure Pending—Number of services that are not configured but are pending reconfiguration. A nonzero value implies that you must initiate a reconfiguration request to provision these services. • Queued—Number of services for which provisioning has not yet been attempted because another service is in the process of being provisioned or deprovisioned. <p>The typical case is when remote device reconfiguration is followed immediately by a new subscriber service provisioning action. In this instance, reconfiguration is allowed to complete, which delays action on the new subscriber service.</p>	extensive
Service Deprovisioning	<p>Status counts for all deprovisioning actions attempted for the remote device.</p> <ul style="list-style-type: none"> • Attempted—Number of service deprovisioning attempts for a service device in the Up state. • Succeeded—Number of services successfully deprovisioned. • Failed—Number of services that failed deprovisioning for all reasons <ul style="list-style-type: none"> • Reject—Number of failures due to an explicit error response during deprovisioning. • Timeout—Number of failures due to no response during deprovisioning. • In Progress—Number of services being deprovisioned but not yet completed. • Queued—Number of services for which deprovisioning has not yet been attempted because another service is in the process of being provisioned or deprovisioned. 	extensive

Table 75: show remote-device-management service-devices Output Fields (*continued*)

Field Name	Field Description	Level of Output
Reconfiguration Requests	Status counts of reconfiguration requests: Received , Succeeded , Failed , and Pending . A reconfiguration request initiated while one is in-progress is rejected.	extensive

Sample Output

show remote-device-management service-devices

```
user@host> show remote-device-management service-devices
```

Device Name	State	Eligible Services	Provisioned Services
Olt-xyz	Up	2	2
Olt-1	Up	3	1
Olt-2	Down	2	0

show remote-device-management service-devices (Device Name)

```
user@host> show remote-device-management service-devices olt-xyz
```

Device Name	State	Eligible Services	Provisioned Services
Olt-xyz	Up	2	2

show remote-device-management service-devices (Device Name Detail)

```
user@host> show remote-device-management service-devices olt-xyz detail
```

```
Device Name: olt-xyz
State: Up
Address: 10.2.3.1
Routing Instance: default
Last State Change Time: Wed Apr 11 07:24:02 2018
Vlan Id List: 1-40 71-80
Dictionary File: /var/home/dict/dictionary-1.xml
Eligible Services: 1
Provisioned Services: 1

Provisioning Method: netconf
  User Name: regress
```



```

Connection Retry Interval: 3
Response Timeout: 10
Response Timeout Count: 3
Bulk Interval: 1000
Bulk Limit: 1
Reconfigure Bulk Limit: 100
Port: 830

```

```

Round Trip Time (millisec)
Minimum      Maximum      Average      Last
1681         1681         1681         1681

```

```

Bulk Count
Minimum      Maximum      Average      Last
1            1            1            1

```

show remote-device-management service-devices (Device Name Extensive)

user@host> show remote-device-management service-devices olt-xyz extensive

```

Device Name: olt-xyz
State: Up
Address: 10.2.3.1
Routing Instance: default
Last State Change Time: Wed Apr 11 07:24:02 2018
Vlan Id List: 1-40 71-80
Dictionary File: /var/home/dict/dictionary-1.xml
Eligible Services: 1
Provisioned Services: 1

```

```

Provisioning Method: netconf
  User Name: regress
  Connection Retry Interval: 3
  Response Timeout: 10
  Response Timeout Count: 3
  Bulk Interval: 1000
  Bulk Limit: 1
  Reconfigure Bulk Limit: 100
  Port: 830

```

```

Round Trip Time (millisec)
Minimum      Maximum      Average      Last
1681         1681         1681         1681

```

Bulk Count				
Minimum	Maximum	Average	Last	
1	1	1	1	
Service Sessions				
Service Session ID	Subscriber Session ID	State	Service Name	
3	2	Active	s1	
Service Provisioning				
Attempted: 1				
Succeeded: 1				
Failed: 0				
Reject: 0				
Timeout: 0				
In Progress: 0				
Reconfigure Pending: 0				
Queued: 0				
Service De-provisioning				
Attempted: 0				
Succeeded: 0				
Failed: 0				
Reject: 0				
Timeout: 0				
In Progress: 0				
Queued: 0				
Reconfiguration Requests				
Received: 0				
Succeeded: 0				
Failed: 0				
Pending: 0				

show remote-device-management services

Syntax

```
show remote-device-management services
session-id id-number
<detail>
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Display information about all service sessions or a specific service session on remote service devices.

Options

none—Display summary information about all remote service sessions.

detail—Display detailed information about remote service sessions.

session-id *id-number*—Identification number for the service session.

Required Privilege Level

view

RELATED DOCUMENTATION

show remote-device-management service-devices		1445
show remote-device-management statistics		1456
show remote-device-management subscribers		1461
show remote-device-management summary		1465

List of Sample Output

- [show remote-device-management services on page 1455](#)
- [show remote-device-management services \(Service Session\) on page 1455](#)
- [show remote-device-management services \(Service Session Detail\) on page 1455](#)

Output Fields

[Table 76 on page 1454](#) lists the output fields for the **show remote-device-management services** command. Output fields are listed in the approximate order in which they appear.

Table 76: show remote-device-management services Output Fields

Field Name	Field Description	Level of Output
Service Session ID	ID number for the remote service session.	All levels
Service State	<p>Status of the service session; represents the RDSM processing state:</p> <ul style="list-style-type: none"> • Active—Remote service session is up and fully provisioned. • Deprovisioning—Remote service session is in the process of being deprovisioned. • Init—Transient state that indicates that the state is transitioning to the provisioning state. • Provisioning—Remote service session is in the process of being provisioned. • Provisioning-Complete—Transient state that indicates service provisioning has completed on the device; transitions to Active or Provisioning-Rollback. • Provisioning-Rollback—Rollback of the service provisioning on the device is in progress. Provisioning is rolled back when provisioning fails for one or more other eligible service devices. 	All levels
Service Name	Name of the remote service.	All levels
Subscriber Session ID	ID number for the subscriber session associated with the service session.	All levels
Service Devices	<p>Detailed status information about each remote device eligible to be provisioned with the named remote service:</p> <ul style="list-style-type: none"> • Name—Name of the service device. • State—State of the remote device service management (RDSM) connection to the service device, Up or Down. • Provisioned—State of the remote service session provisioning: <ul style="list-style-type: none"> • Yes—Service has been successfully provisioned for the service device. • No—Service provisioning is in progress or has not been performed. Provisioning might not have been performed because the remote device transitioned from Down to Up and requires reconfiguration. When a remote device is declared Down, it is no longer provisioned and must be reconfigured. 	detail

Sample Output

show remote-device-management services

```
user@host> show remote-device-management services
```

Service Session ID	Subscriber Session ID	State	Service Name
1234	111	Active	s1
222	45	Provisioning	s1
555	100	Deprovisioning	s2

show remote-device-management services (Service Session)

```
user@host> show remote-device-management services session-id 1234
```

Service Session ID	Subscriber Session ID	State	Service Name
1234	111	Active	s1

show remote-device-management services (Service Session Detail)

```
user@host> show remote-device-management services session-id 1234 detail
```

```
Service Session
Service Session ID: 1234
Service State: Active
Service Name: SUBSCRIBER-UPSTREAM-PROFILE
Subscriber Session ID: 111

Service Devices
Device Name      State      Provisioned
olt-xyz          Up         Yes
olt-1            Up         No
olt-3            Down       No
```

show remote-device-management statistics

Syntax

```
show remote-device-management statistics
(summary | service-devices device-name)
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Display a global summary of service statistics for all remote devices or detailed statistics for a specific remote service device.

Options

service-devices *device-name*—Display statistics for the specified service device.

summary—Display a summary of service statistics for all remote devices.

Required Privilege Level

view

RELATED DOCUMENTATION

show remote-device-management service-devices 1445
show remote-device-management services 1453
show remote-device-management subscribers 1461
show remote-device-management summary 1465

List of Sample Output

[show remote-device-management statistics \(Summary\) on page 1458](#)

[show remote-device-management statistics \(Service Devices\) on page 1459](#)

Output Fields

[Table 77 on page 1457](#) lists the output fields for the **show remote-device-management statistics** command. Output fields are listed in the approximate order in which they appear.

Table 77: show remote-device-management statistics Output Fields

Field Name	Field Description
Service Activations	<p>Status of service profile instantiation requests from authd to provision remote services. Each service activation results in a service provisioning request to each remote device. The values displayed are the totals for all remote devices.</p> <ul style="list-style-type: none"> • Received—Number of service provisioning requests received by RDSM. • Acked—Number of service provisioning requests acknowledged by RDSM. • Nacked—Number of service provisioning requests not acknowledged by RDSM. • In progress—Number of service provisioning requests currently in progress on remote devices.
Service Deactivations	<p>Status of service profile deinstantiation request from authd to provision remote services. Each service deactivation results in a service deprovisioning request to each remote device. The values displayed are the totals for all remote devices.</p> <ul style="list-style-type: none"> • Received—Number of service deprovisioning requests received by RDSM. • Acked—Number of service deprovisioning requests acknowledged by RDSM. • Nacked—Number of service deprovisioning requests not acknowledged by RDSM. • In progress—Number of service deprovisioning requests currently in progress on remote devices.
Round Trip Time (millisec)	<p>Aggregate duration for all dictionary RPCs required to provision or deprovision a service: Minimum, Maximum, Average, and Last.</p> <p>When the bulk limit is set to 1, these statistics are reported per remote service.</p> <p>When the bulk limit is greater than 1, the duration includes all services that are part of the bulk configuration.</p>
Bulk Count	<p>Number of services included in a bulk configuration: Minimum, Maximum, Average, and Last. The number is determined by the bulk-limit and bulk-interval configuration for the remote device at the <code>[edit system services remote-device-management service-device <i>device-name</i> provisioning-method netconf]</code> hierarchy level.</p> <p>When the bulk limit is set to 1, the count is 1 for all fields.</p>

Table 77: show remote-device-management statistics Output Fields (*continued*)

Field Name	Field Description
Service Provisioning	<p>Status counts for all provisioning actions attempted for the remote device.</p> <ul style="list-style-type: none"> • Attempted—Number of service provisioning attempts for a service device in the Up state. • Succeeded—Number of services successfully provisioned. • Failed—Number of services that failed provisioning for all reasons <ul style="list-style-type: none"> • Reject—Number of failures due to an explicit error response during provisioning. • Timeout—Number of failures due to no response during provisioning. • In Progress—Number of services being provisioned but not yet completed. • Reconfigure Pending—Number of services that are not configured but are pending reconfiguration. A nonzero value implies that you must initiate a reconfiguration request to provision these services. • Queued—Number of services for which provisioning has not yet been attempted because another service is in the process of being provisioned or deprovisioned. <p>The typical case is when remote device reconfiguration is followed immediately by a new subscriber service provisioning action. In this instance, reconfiguration is allowed to complete, which delays action on the new subscriber service.</p>
Service Deprovisioning	<p>Status counts for all deprovisioning actions attempted for the remote device.</p> <ul style="list-style-type: none"> • Attempted—Number of service deprovisioning attempts for a service device in the Up state. • Succeeded—Number of services successfully deprovisioned. • Failed—Number of services that failed deprovisioning for all reasons <ul style="list-style-type: none"> • Reject—Number of failures due to an explicit error response during deprovisioning. • Timeout—Number of failures due to no response during deprovisioning. • In Progress—Number of services being deprovisioned but not yet completed. • Queued—Number of services for which deprovisioning has not yet been attempted because another service is in the process of being provisioned or deprovisioned.
Reconfiguration Requests	<p>Status counts of reconfiguration requests: Received, Succeeded, Failed, and Pending. A reconfiguration request initiated while one is in-progress is rejected.</p>

Sample Output

show remote-device-management statistics (Summary)

```
user@host> show remote-device-management statistics summary
```


Service Activations

Received: 10
 Acked: 7
 Nackd: 2
 In Progress: 1

Service Deactivations

Received: 4
 Acked: 3
 Nackd: 1
 In Progress: 0

show remote-device-management statistics (Service Devices)

user@host> **show remote-device-management statistics service-devices olt-xyz**

Round Trip Time (millisec)

Minimum	Maximum	Average	Last
1000	2000	1300	1500

Bulk Count

Minimum	Maximum	Average	Last
1	5	2	3

Service Provisioning

Attempted: 2
 Succeeded: 1
 Failed: 0
 Reject: 0
 Timeout: 0
 In Progress: 1
 Reconfigure Pending: 1
 Queued: 0

Service Deprovisioning

Attempted: 0
 Succeeded: 0
 Failed: 0
 Reject: 0
 Timeout: 0
 In Progress: 0
 Queued: 0

Reconfiguration Requests

Received: 2

Succeeded: 1

Failed: 0

Pending: 1

show remote-device-management subscribers

Syntax

```
show remote-device-management subscribers
session-id id-number
<detail>
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Display information about service sessions for all subscriber sessions or about all service sessions for a specific subscriber session on remote service devices.

Options

none—Display summary information about remote service sessions for all subscribers.

detail—Display detailed information about remote service sessions for all subscribers.

session-id *id-number*—Identification number for the subscriber session.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear remote-device-management statistics | 1267](#)

[show remote-device-management service-devices | 1445](#)

[show remote-device-management services | 1453](#)

[show remote-device-management statistics | 1456](#)

[show remote-device-management summary | 1465](#)

List of Sample Output

[show remote-device-management subscribers on page 1463](#)

[show remote-device-management subscribers \(Subscriber Session\) on page 1463](#)

[show remote-device-management subscribers \(Subscriber Session Detail\) on page 1463](#)

Output Fields

[Table 78 on page 1462](#) lists the output fields for the **show remote-device-management subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 78: show remote-device-management subscribers Output Fields

Field Name	Field Description	Level of Output
Subscriber Session ID	ID number for the subscriber session associated with the service session.	All levels
Service Session ID	ID number for the remote service session.	All levels
State	<p>Status of the service session; represents the RDSM processing state:</p> <ul style="list-style-type: none"> • Active—Remote service session is up and fully provisioned. • Deprovisioning—Remote service session is in the process of being deprovisioned. • Init—Transient state that indicates that the state is transitioning to the provisioning state. • Provisioning—Remote service session is in the process of being provisioned. • Provisioning-Complete—Transient state that indicates service provisioning has completed on the device; transitions to Active or Provisioning-Rollback. • Provisioning-Rollback—Rollback of the service provisioning on the device is in progress. Provisioning is rolled back when provisioning fails for one or more other eligible service devices. 	All levels
Service Name	Name of the remote service.	All levels
Service Devices	<p>Detailed status information about each remote device eligible to be provisioned with the named remote service:</p> <ul style="list-style-type: none"> • Name—Name of the service device. • State—State of the remote device service management (RDSM) connection to the service device, Up or Down. • Provisioned—State of the remote service session provisioning: <ul style="list-style-type: none"> • Yes—Service has been successfully provisioned for the service device. • No—Service provisioning is in progress or has not been performed. Provisioning might not have been performed because the remote device transitioned from Down to Up and requires reconfiguration. When a remote device is declared Down, it is no longer provisioned and must be reconfigured. 	detail

Sample Output

show remote-device-management subscribers

```
user@host> show remote-device-management subscribers
```

Subscriber Session ID	Service Session ID	State	Service Name
111	1234	Active	s1
1238	111	Active	s112
222	45	Provisioning	s1
555	100	Deprovisioning	s2

show remote-device-management subscribers (Subscriber Session)

```
user@host> show remote-device-management subscribers session-id 111
```

Subscriber Session ID	Service Session ID	State	Service Name
111	1234	Active	s1
111	1238	Active	s112

show remote-device-management subscribers (Subscriber Session Detail)

```
user@host> show remote-device-management subscribers session-id 111 detail
```

Subscriber Session ID: 111

Service Session

Service Session ID: 1234

Service State: Active

Service Name: s1

Subscriber Session ID: 111

Service Devices

Device Name	State	Provisioned
olt-xyz	Up	Yes
olt-1	Up	No
olt-3	Down	No

Service Session

Service Session ID: 1238

Service State: Active

Service Name: s112

Subscriber Session ID: 111

Service Devices		
Device Name	State	Provisioned
olt-xyz	Up	Yes
olt-1	Up	No
olt-3	Down	No

show remote-device-management summary

Syntax

```
show remote-device-management summary;
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Display summary information about the remote service devices, such as session state and service state.

Required Privilege Level

view

RELATED DOCUMENTATION

- [show remote-device-management service-devices | 1445](#)
- [show remote-device-management services | 1453](#)
- [show remote-device-management statistics | 1456](#)
- [show remote-device-management subscribers | 1461](#)

List of Sample Output

[show remote-device-management summary on page 1467](#)

Output Fields

[Table 79 on page 1465](#) lists the output fields for the **show remote-device-management summary** command. Output fields are listed in the approximate order in which they appear.

Table 79: show remote-device-management summary Output Fields

Field Name	Field Description
Service Devices by Connection State	Number of services devices where the remote device service management (RDSM) connection to the remote device is in the Up or Down state.

Table 79: show remote-device-management summary Output Fields (*continued*)

Field Name	Field Description
Service Sessions by State	<p>Number of service sessions in each of the following states, which represent the RDSM processing state:</p> <ul style="list-style-type: none"> • Active—Remote service session is up and fully provisioned. • Deprovisioning—Remote service session is in the process of being deprovisioned. • Init—Transient state that indicates that the state is transitioning to the provisioning state. • Provisioning—Remote service session is in the process of being provisioned. • Provisioning-Complete—Transient state that indicates service provisioning has completed on the device; transitions to Active or Provisioning-Rollback. • Provisioning-Rollback—Rollback of the service provisioning on the device is in progress. Provisioning is rolled back when provisioning fails for one or more other eligible service devices.
Service Activations	<p>Status of service profile instantiation requests from authd to provision remote services. Each service activation results in a service provisioning request to each remote device. The values displayed are the totals for all remote devices.</p> <ul style="list-style-type: none"> • Received—Number of service provisioning requests received by remote devices. • Acked—Number of service provisioning requests acknowledged by remote devices. • Nacked—Number of service provisioning requests not acknowledged by remote devices. • In progress—Number of service provisioning requests currently in progress on remote devices.

Table 79: show remote-device-management summary Output Fields (*continued*)

Field Name	Field Description
Service Deactivations	<p>Status of service profile deinstantiation request from authd to provision remote services. Each service deactivation results in a service deprovisioning request to each remote device. The values displayed are the totals for all remote devices.</p> <ul style="list-style-type: none"> • Received—Number of service deprovisioning requests received by remote devices. • Acked—Number of service deprovisioning requests acknowledged by remote devices. • Nacked—Number of service deprovisioning requests not acknowledged by remote devices. • In progress—Number of service deprovisioning requests currently in progress on remote devices.

Sample Output

show remote-device-management summary

user@host> **show remote-device-management summary**

Service Devices by Connection State

Up: 2
Down: 1

Service Sessions by State

Active: 2
Provisioning: 1
Deprovisioning: 0

Service Activations

Received: 10
Acked: 7
Nacked: 2
In progress: 1

Service De-activations

Received: 4
Acked: 3

Nacked: 1
In progress: 0

show services application-identification application

Syntax

```
show services application-identification application
<detail <application-name> | summary >
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display detailed information about a specified application signature, all application signatures, or a summary of the existing application signatures and nested application signatures. Both custom and predefined application signatures and nested application signatures can be displayed.

Options

none—(Same as summary) Display a summary of the application identification application information.

detail <application-name> | summary—(Optional) Display the specified level of output.

application-name—(Optional) Display detailed information for the specified application name; maximum 31 characters. Predefined applications have the prefix **junos-** to avoid conflict with user-defined ones.

Required Privilege Level

view

RELATED DOCUMENTATION

[request services application-identification install | 1300](#)

[request services application-identification application | 1294](#)

List of Sample Output

[show services application-identification application summary on page 1471](#)

[show services application-identification application detail on page 1471](#)

[show services application-identification application detail \(Specific Application\) on page 1475](#)

[show services application-identification application detail \(Specific Application\) on page 1476](#)

Output Fields

[Table 80 on page 1470](#) lists the output fields for the **show services application-identification application** command. Output fields are listed in the approximate order in which they appear.

Table 80: show services application-identification application Output Fields

Field Name	Field Description	Level of Output
Application(s)	Number of applications present.	none summary
Application	Name of the predefined application.	none summary
Disabled	Status (Yes or No) of the application and whether the mapping method is currently used to identify this application.	none summary
Application ID	Unique ID number of an application. ID numbers 1 through 32,767 are automatically generated for predefined applications; these IDs do not change.	none summary
Order	Unique number used to specify priority when multiple applications match the traffic. The lowest order number takes the highest priority. The order attribute is applicable only for custom signatures.	none summary
Application Name	Name of the predefined application.	detail
Application type	Basic application type, such as HTTP.	detail
Description	Description of the predefined application.	detail
Number of Parent Group(s)	Number of parent groups associated with this application.	detail
Application Tags	<p>Category specifying one or more following attributes of the application:</p> <p>characteristic: One or more characteristics of the application.</p> <p>risk: Level of risk of the application.</p> <p>subcategory: Subcategory of the application.</p> <p>category: Technology of the application.</p>	detail

Table 80: show services application-identification application Output Fields (continued)

Field Name	Field Description	Level of Output
Layer-7 Protocol(s)	Layer 7 protocols associated with the application.	detail
Port Mapping Default port	Ports associated with the application.	detail
Signature	Signature mapping criteria for application identification: Port range , Client-to-server , and Order .	detail

Sample Output

show services application-identification application summary

user@host> **show services application-identification application summary**

```

Application(s): 2564
  Applications                               Disabled      ID      Order
  junos:DOT-NET                             No            10182   2564

  junos:ICMP-PHOTURIS-NEED-AUTHOR            No            11377   2563

  junos:MYSPACE-TAG-ME                       No            10683   2562

  junos:SLACKER                              No            1179    2561

  junos:ICMP-TYPE-55                         No            11392   2560

  junos:FLIPDRIVE-SSL                        No            10939   2559

  junos:ICMP-MOBILE-HOST-REDIR               No            11363   2558

  junos:TWITPIC                              No            864     2557

  junos:ICMP-TYPE-245                        No            11582   2556

```

show services application-identification application detail

user@host> **show services application-identification application detail**

re0:

```

-----
Application Name: junos:dot-net
Application type: DOT-NET
Description: .Net Remoting
Application ID: 10182
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:infrastructure:rpc
Port Mapping:
    Default ports: N/A
Signature:
    Port range: N/A
    Client-to-server
    Order: 1
Application Name: junos:icmp-photuris-need-author
Application type: ICMP-PHOTURIS-NEED-AUTHOR
Description: ICMP Type 40 Code 5 - Photuris (Need Authorization)
Application ID: 11377
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:infrastructure:networking
Port Mapping:
    Default ports: N/A
Signature:
    Port range: N/A
    Client-to-server
    Order: 5
Application Name: junos:myspace-tag-me
Application type: MYSPACE-TAG-ME
Description: This signature detects Tag Me by BitRhymes on MySpace Apps.  Tag
             Me by BitRhymes on MySpace Apps is a Web-based entertainment
             application on the popular social network MySpace.
Application ID: 10683
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:web:social-networking
Port Mapping:
    Default ports: N/A
Signature:
    Port range: N/A

```

```

    Client-to-server
    Order: 4
Application Name: junos:slacker
Application type: SLACKER
Description: This protocol plug-in classifies the http traffic to the host
             .slackr.com.
Application ID: 1179
Disabled: No
Number of Parent Group(s): 2
Application Groups:
    junos:multimedia:divers
    junos:multimedia
Port Mapping:
    Default ports: N/A
Signature:
    Port range: N/A
    Client-to-server
    Order: 3
Application Name: junos:icmp-type-55
Application type: ICMP-TYPE-55
Description: ICMP Type 55 - Unassigned
Application ID: 11392
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:infrastructure:networking
Port Mapping:
    Default ports: N/A
Signature:
    Port range: N/A
    Client-to-server
    Order: 2
Application Name: junos:flipdrive-ssl
Application type: FLIPDRIVE-SSL
Description: This signature detects logins to FlipDrive, a cloud-based
             file-sharing and backup service.
Application ID: 10939
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:web:file-sharing
Port Mapping:
    Default ports: N/A
Signature:

```

```

    Port range: N/A
    Client-to-server
    Order: 1
Application Name: junos:icmp-mobile-host-redir
Application type: ICMP-MOBILE-HOST-REDIR
Description: ICMP Type 32 - Mobile Host Redirect
Application ID: 11363
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:infrastructure:networking
Port Mapping:
    Default ports: N/A
Signature:
    Port range: N/A
    Client-to-server
    Order: 5
Application Name: junos:twitpic
Application type: TWITPIC
Description: This signature detects Twitpic, a Web site that allows users to
            easily post pictures to the Twitter microblogging and social media
            service.
Application ID: 864
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:web:social-networking
Port Mapping:
    Default ports: N/A
Signature:
    Port range: N/A
    Client-to-server
    Order: 4
Application Name: junos:icmp-type-245
Application type: ICMP-TYPE-245
Description: ICMP Type 245 - Unassigned
Application ID: 11582
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:infrastructure:networking
Port Mapping:
    Default ports: N/A
Signature:

```



```

Port range: N/A
Client-to-server
Order: 3

---(more)---
```

show services application-identification application detail (Specific Application)

user@host> show services application-identification application detail junos:SKYPE

```

Application Name: junos:SKYPE
Application type: SKYPE
Description: This signature detects Skype, which is a proprietary P2P VOIP
              network. It is a "complete black box" for both users and
              analyzers. It uses security through obscurity to make itself
              troublesome to analyze or reverse-engineer without a significant
              amount of work, or use of emulation. It uses AES block cipher, the
              RSA public key cryptosystem, the ISO 9796-2 signature padding
              scheme, the SHA-1 hash function, and the RC4 stream cipher through
              the communications between the client to client, client to
              supernodes and supernode to supernode.

Application ID: 183
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:web:infrastructure:voip
Application Tags:
    characteristic      : Supports File Transfer
    characteristic      : Evasive
    characteristic      : Bandwidth Consumer
    risk                 : 4
    subcategory          : VOIP
    category             : Infrastructure
Layer-7 Protocol(s):  UDP      / 216
                      TCP      / 205
                      SSL      / 199
                      HTTPS    / 68
                      HTTP     / 67

Port Mapping:
    Default ports: N/A
Signature:
    Port range: N/A
    Client-to-server
    Order: 20
```

show services application-identification application detail (Specific Application)

user@host> **show services application-identification detail junos:http**

```
re0:
-----
Application Name: junos:http
Application type: HTTP
Description: This signature detects HyperText Transfer Protocol (HTTP), which
             is a protocol used by the World Wide Web. It defines how messages
             are formatted and transmitted and what actions Web servers and
             browsers should take in response to various commands. HTTP usually
             runs on TCP port 80.
Application ID: 67
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:web
Port Mapping:
    Default ports: TCP/80,3128,8000,8080
Signature:
    Port range: N/A
    Client-to-server
    Order: 3
```

show services application-identification application-system-cache

Syntax

```
show services application-identification application-system-cache  
<interface interface-name>
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display the database of cached values stored by the application identification system.

NOTE: The **show services application-identification application-system-cache** command gives the information only when the application identifier (AI) is matched with the signature.

Options

none—Display the database of cached values for the all services interfaces.

interface *interface-name*—(Optional) Display the database of cached values for the specified services interface.

Required Privilege Level

view

RELATED DOCUMENTATION

[request services application-identification application](#) | 1294

List of Sample Output

[show services application-identification application-system-cache on page 1480](#)

[show services application-identification application-system-cache interface on page 1480](#)

Output Fields

[Table 81 on page 1478](#) lists the output fields for the **show services application-identification application-system-cache** command. Output fields are listed in the approximate order in which they appear.

Table 81: show services application-identification application-system-cache Output Fields

Field Name	Field Description
application-cache	Status (on or off) of the application cache.
cache-entry-timeout	Number of seconds the mapping information is saved.
pic	PIC number of the accumulated statistics.
IP address	IP address of the traffic flow for which application-identification is enabled.
Port	Port number of the traffic flow for which application-identification is enabled.
Protocol	Protocol name of the flow for which application-identification is enabled.
Application	Application number, which is a unique identifier that denotes the application or service for which identification of traffic flows is enabled.
Classification Path	Protocols or nested applications that denote the paths traversed for classified packets.
PIC	PIC number of the accumulated statistics. For the interface on which deep packet inspection (DPI) application is not running, that detail is also displayed for the corresponding interface.
Unknown applications	Number of unknown applications.
Cache hits	Number of sessions that matched the application in the application identification cache.
Cache misses	Number of sessions that did not find the application in the application identification cache.
Client-to-server packets processed	Number of client-to-server packets processed.
Server-to-client packets processed	Number of server-to-client packets processed.
Client-to-server bytes processed	Number of client-to-server payload bytes processed.

Table 81: show services application-identification application-system-cache Output Fields (*continued*)

Field Name	Field Description
Server-to-client layer bytes processed	Number of server-to-client payload bytes processed.
Client-to-server packets processed	Number of client-to-server packets processed.
Server-to-client packets processed	Number of server-to-client packets processed.
Client-to-server bytes processed	Number of client-to-server payload bytes processed.
Server-to-client layer bytes processed	Number of server-to-client payload bytes processed.
Sessions bypassed due to resource allocation failure	Number of sessions bypassed due to resource allocation failure.
Segment case 1 - New segment to left	Number of TCP segments contained before the previous segment.
Segment case 2 - New segment overlap right	Number of TCP segments that start before the previous segment and are contained in it.
Segment case 3 - Old segment overlapped	Number of TCP segments that start before the previous segment and extend beyond it.
Segment case 4 - New segment overlapped	Number of TCP segments that start and end within the previous segment.

Table 81: show services application-identification application-system-cache Output Fields (*continued*)

Field Name	Field Description
Segment case 5 - New segment overlap left	Number of TCP segments that start within the previous segments and extend beyond it.
Segment case 6 - New segment to right	Number of TCP segments that start after the previous segment. This is the normal case.

Sample Output

show services application-identification application-system-cache

```
user@host> show services application-identification application-system-cache
```

```
Application System Cache Configurations:
  application-cache: on
  cache-entry-timeout: 3600 seconds
pic: ams0
pic: ms-0/3/0
ms-0/3/0 is not running DPI engine
pic: ams1
pic: ms-0/0/0
IP address: 192.0.2.2                      Port: 80      Protocol: TCP
Application: HTTP:YOUTUBE
Classification Path: IP:TCP:HTTP:YOUTUBE
```

show services application-identification application-system-cache interface

```
user@host> show services application-identification application-system-cache interface ms-1/0/0
```

```
Application System Cache Configurations:
  application-cache: on
  cache-entry-timeout: 3600 seconds
pic: ms-0/0/0
IP address: 192.0.2.2                      Port: 80      Protocol: TCP
Application: HTTP:YOUTUBE
Classification Path: IP:TCP:HTTP:YOUTUBE
user@host> show services application-identification counter
```

pic: ams0

ms-0/3/0 is not running DPI engine

pic: ams1

Counter type	Value
Unknown applications	32682
Cache hits	323504
Cache misses	400
Client-to-server packets processed	2034
Server-to-client packets processed	1982
Client-to-server bytes processed	258786
Server-to-client bytes processed	1314722
Sessions bypassed due to resource allocation failure	0
Segment case 1 - New segment to left	0
Segment case 2 - New segment overlap right	0
Segment case 3 - Old segment overlapped	0
Segment case 4 - New segment overlapped	0
Segment case 5 - New segment overlap left	0
Segment case 6 - New segment to right	0

show services application-identification commit-status (Next Gen Services)

Syntax

```
show services application-identification commit-status]
```

Release Information

Statement introduced in Junos OS Release 19.3R1 on MX Series running Next Gen Services.

Description

Display information about the commit status. Because the custom signatures commit is performed asynchronously, the command output shows the current status of your configuration commit.

Required Privilege Level

view

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

List of Sample Output

[show services application-identification commit-status on page 1482](#)

[show services application-identification commit-status on page 1482](#)

[show services application-identification commit-status on page 1483](#)

Sample Output

show services application-identification commit-status

```
user@host> show services application-identification commit-status
```

```
Custom signatures commit is in progress
```

show services application-identification commit-status

```
user@host> show services application-identification commit-status
```



```
Custom signatures committed successfully
```

show services application-identification commit-status

```
user@host> show services application-identification commit-status
```

```
Custom signatures serialization failed
```

show services application-identification counter

Syntax

```
show services application-identification counter
<interface interface-name>
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display application identification counter statistics.

Options

none—Display counter statistics for all services interfaces.

interface *interface-name*—(Optional) Display counter statistics for the specified services interface.

Required Privilege Level

view

RELATED DOCUMENTATION

- [Application Identification Overview | 431](#)
- [Configuring Custom Application Signatures | 435](#)

List of Sample Output

[show services application-identification counter on page 1485](#)

Output Fields

[Table 82 on page 1484](#) lists the output fields for the **show services application-identification counter** command. Output fields are listed in an approximate order in which they appear.

Table 82: show services application-identification counter Output Fields

Field Name	Field Description
PIC	PIC number of the accumulated statistics.
Unknown applications	Number of unknown applications.
Cache hits	Number of sessions that matched the application in the application identification cache.

Table 82: show services application-identification counter Output Fields (*continued*)

Field Name	Field Description
Cache misses	Number of sessions that did not find the application in the application identification cache.
Client-to-server packets processed	Number of client-to-server packets processed.
Server-to-client packets processed	Number of server-to-client packets processed.
Client-to-server bytes processed	Number of client-to-server payload bytes processed.
Server-to-client layer bytes processed	Number of server-to-client payload bytes processed.
Sessions bypassed due to resource allocation failure	Number of sessions bypassed due to resource allocation failure.
Segment case 1 - New segment to left	Number of TCP segments contained before the previous segment.
Segment case 2 - New segment overlap right	Number of TCP segments that start before the previous segment and are contained in it.
Segment case 3 - Old segment overlapped	Number of TCP segments that start before the previous segment and extend beyond it.
Segment case 4 - New segment overlapped	Number of TCP segments that start and end within the previous segment.
Segment case 5 - New segment overlap left	Number of TCP segments that start within the previous segments and extend beyond it.
Segment case 6 - New segment to right	Number of TCP segments that start after the previous segment. This is the normal case.

Sample Output

```
show services application-identification counter
```

```
user@host> show services application-identification counter
```

pic: 5/0

Counter type	Value
Unknown applications	0
Cache hits	0
Cache misses	36
Client-to-server packets processed	16
Server-to-client packets processed	101
Client-to-server bytes processed	3494
Server-to-client bytes processed	112493
Sessions bypassed due to resource allocation failure	0
Segment case 1 - New segment to left	11
Segment case 2 - New segment overlap right	8
Segment case 3 - Old segment overlapped	0
Segment case 4 - New segment overlapped	0
Segment case 5 - New segment overlap left	0
Segment case 6 - New segment to right	7

pic: 5/1

Counter type	Value
Unknown applications	0
Cache hits	0
Cache misses	0
Client-to-server packets processed	0
Server-to-client packets processed	0
Client-to-server bytes processed	0
Server-to-client bytes processed	0
Sessions bypassed due to resource allocation failure	0
Segment case 1 - New segment to left	0
Segment case 2 - New segment overlap right	0
Segment case 3 - Old segment overlapped	0
Segment case 4 - New segment overlapped	0
Segment case 5 - New segment overlap left	0
Segment case 6 - New segment to right	0

show services application-identification group

Syntax

```
show services application-identification group [detail application-group name | summary]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display detailed or summary information about a specified application signature group or all application signature groups. Both custom and predefined application signature groups can be displayed.

Options

none—Display summary information for all application signature groups.

detail | summary—Display the specified level of output.

application-name—Application name; maximum 31 characters. Predefined applications have the prefix **junos-** to avoid conflict with user-defined ones.

Required Privilege Level

view

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

[request services application-identification group | 1298](#)

List of Sample Output

[show services application-identification group summary on page 1489](#)

[show services application-identification group detail on page 1490](#)

Output Fields

[Table 83 on page 1488](#) lists the output fields for the **show services application-identification group** command. Output fields are listed in the approximate order in which they appear.

Table 83: show services application-identification group Output Fields

Field Name	Field Description	Level of Output
Group ID	Unique ID number of an application signature group. ID numbers 1 through 32,767 are automatically generated for predefined application signatures and application signature groups; these IDs do not change. ID numbers for custom application signatures and application signature groups use ID numbers 32,768 through 65,534.	none detail summary
Disabled	Status of the application signature group and whether the signature method is currently used to identify this application. The default is No.	none summary
Application Group(s)	Number of application signature groups present.	none summary
Applications	Names of application signatures associated with this application signature group.	none detail summary
Group Name	Name of an application signature or application signature group.	detail
Description	Description of the specified application in the detailed display. If a description is not previously specified, N/A is displayed for this field.	detail
Number of Applications	Total number of applications contained in the group.	detail
Number of Sub-Groups	Total number of sub-groups associated with this application signature group.	detail
Number of Parent-Groups	Total number of parent groups in this application signature group or cluster.	detail
Sub-Group(s)	Application signature sub-groups present.	detail

Sample Output

show services application-identification group summary

user@host> **show services application-identification group summary**

```
Application Group(s): 66
Application Groups           Disabled  ID
junos:web:social-networking:facebook    No      68
junos:web:reference                     No      67
junos:infrastructure:legacy              No      66
junos:web:cdn                           No      65
junos:infrastructure:scada               No      64
junos:web:real-estate                    No      63
junos:web:finance                        No      62
junos:multimedia:audio-streaming         No      61
junos:web:remote-access                  No      60
junos:web:p2p                           No      59
junos:remote-access:backdoors            No      58
junos:infrastructure:authentication      No      57
junos:web:forums                         No      56
junos:remote-access:command              No      55
junos:infrastructure:scm                 No      54
junos:web:portal                         No      53
junos:web:shopping                       No      52
junos:infrastructure:rpc                 No      51
junos:messaging:mail                    No      50
junos:web:search                         No      49
junos:infrastructure:encryption          No      48
junos:gaming:divers                      No      47
junos:p2p:file-sharing                   No      46
junos:infrastructure:backup              No      45
junos:multimedia:transport               No      44
junos:gaming:protocols                   No      43
junos:web:advertisements                 No      42
junos:infrastructure:monitoring          No      41
junos:infrastructure:mobile              No      40
junos:infrastructure:file-servers        No      39
junos:web:infrastructure                 No      38
junos:web:wiki                           No      37
junos:web:image-sharing                  No      36
junos:infrastructure:directory           No      35
junos:infrastructure:database            No      34
junos:remote-access:tunneling            No      33
junos:remote-access:interactive-desktop  No      32
```

junos:web:gaming	No	31
junos:web:anonymizer	No	30
junos:web:blogging	No	29
junos:remote-access:divers	No	28
junos:remote-access	No	27
junos:p2p:divers	No	26
junos:p2p	No	25
junos:web:news	No	24
junos:gaming:web-based	No	23
junos:gaming	No	22
junos:web:messaging	No	21
junos:multimedia:web-based	No	20
junos:web:file-sharing	No	19
junos:web:travel	No	18
junos:multimedia:video-streaming	No	17
junos:messaging:instant-messaging	No	16
junos:web:multimedia	No	15
junos:infrastructure:voip	No	14
junos:messaging:divers	No	13
junos:messaging	No	12
junos:web:applications	No	11
junos:multimedia:divers	No	10
junos:multimedia	No	9
junos:web:divers	No	8
junos:web:social-networking	No	7
junos:web	No	6
junos:infrastructure:networking	No	5
junos:infrastructure:divers	No	4
junos:infrastructure	No	3

show services application-identification group detail

user@host> **show services application-identification group detail junos:social-networking**

```

Group Name: junos:web
Group ID: 15
Description: N/A
Disabled: No
Number of Applications: 1
Number of Sub-Groups: 21
Number of Parent-Groups: 1
Applications:
    junos:http
Sub Groups:

```



```
junos:web:forums  
junos:web:travel  
junos:web:reference  
junos:web:portal  
junos:web:blogging  
junos:web:shopping  
junos:web:search  
junos:web:anonymizer  
junos:web:image-sharing  
junos:web:file-sharing  
junos:web:remote-access  
junos:web:real-estate  
junos:web:news  
junos:web:gaming  
junos:web:p2p  
junos:web:applications  
junos:web:multimedia  
junos:web:divers  
junos:web:messaging  
junos:web:social-networking  
junos:web:infrastructure
```

show services application-identification statistics application-groups

Syntax

```
show services application-identification statistics application-groups
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display cumulative session and byte statistics per application group. Statistics are displayed in alphabetical order.

Options

This command has no options.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear services application-identification statistics](#) | [1270](#)

List of Sample Output

[show services application-identification statistics application-groups](#) on [page 1493](#)

Output Fields

[Table 84 on page 1492](#) lists the output fields for the **show services application-identification statistics application-groups** command. Output fields are listed in the approximate order in which they appear.

Table 84: show services application-identification statistics application-groups Output Fields

Field Name	Field Description
Last Reset	Date, time, and how long ago the statistics for the sessions were cleared. The format None specified is <i>year-month-day hour:minute:second timezone</i> . If you did not clear the statistics previously at any point, Never is displayed.
Application Group	Name of the application group.
Sessions	Number of sessions for the application group.
Kilo Bytes	Size of the application group in kilobytes.

Sample Output

show services application-identification statistics application-groups

user@host> **show services application-identification statistics application-groups**

Last Reset: 2014-02-19 00:38:01 PST

Application Group	Sessions	Kilo Bytes
junos:infrastructure	2	18
junos:infrastructure:monitoring	2	18

show services application-identification statistics applications

Syntax

```
show services application-identification statistics applications
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display cumulative session and byte statistics per application. Statistics are displayed in alphabetical order.

Options

This command has no options.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear services application-identification statistics](#) | [1270](#)

List of Sample Output

[show services application-identification statistics applications on page 1495](#)

Output Fields

[Table 85 on page 1494](#) lists the output fields for the **show services application-identification statistics applications** command. Output fields are listed in the approximate order in which they appear.

Table 85: show services application-identification statistics applications Output Fields

Field Name	Field Description
Last Reset	Date, time, and how long ago the statistics for the sessions were cleared in the format <i>year-month-day hour:minute:second timezone</i> . If you did not clear the statistics previously at any point, Never is displayed.
Application	Name of the application.
Sessions	Number of sessions for the application.
Bytes	Size of the application in bytes.

Sample Output

show services application-identification statistics applications

user@host> **show services application-identification statistics applications**

Last Reset: 2014-01-26 18:32:36 PST

Application	Sessions	Bytes
junos:http	4	24009
junos:https	1	101823
junos:hulu	1	48329
junos:linkedin	1	2650
junos:netflix	2	32747

show services application-identification status

Syntax

```
show services application-identification status
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display detailed information about application identification status.

Required Privilege Level

view

RELATED DOCUMENTATION

[Application Identification Overview | 431](#)

[Configuring Custom Application Signatures | 435](#)

[request services application-identification application | 1294](#)

List of Sample Output

[show services application-identification status on page 1497](#)

Output Fields

[Table 86 on page 1496](#) lists the output fields for the **show services application-identification status** command. Output fields are listed in the approximate order in which they appear.

Table 86: show services application-identification status Output Fields

Field Name	Field Description
Application Identification	Details of the application-identification engine and the processing details of sessions.
Status	Status of application identification: Enabled or Disabled .
Sessions under app detection	Number of sessions undergoing application identification detection.
Engine Version	Application identification detector engine version.
Max TCP session packet memory	Maximum number of TCP sessions that application identification maintains.

Table 86: show services application-identification status Output Fields *(continued)*

Field Name	Field Description
Force packet plugin	Force packet plugin status: Enabled or Disabled .
Force stream plugin	Force stream plugin status: Enabled or Disabled .
Statistics collection interval	Frequency (in minutes) for collecting statistics.
Application System Cache	Details of entries in the application system cache.
Status	Status of application system cache: Enabled or Disabled .
Max Number of entries in cache	Maximum number of cache entries.
Cache timeout	Number of seconds after which the cache entries expires.
Protocol Bundle	Information regarding application package downloads.
Download Server CGI	URL of the server from where protocol bundle was downloaded.
Auto Update	Status of auto update to receive protocol bundle updates from the server: Enabled or Disabled .
Slot	Number of the slot pertaining to the packets for which application-identification is associated.
Status	Status of protocol bundle: Active or Free .
Version	Version of protocol bundle.
Session	Number of active sessions.

Sample Output

show services application-identification status

user@host> show services application-identification status

```
pic: 5/0
```

```
Application Identification
```

Status	Enabled
Sessions under app detection	0
Engine Version	4.18.1-20 (build date Feb 15 2014)
Max TCP session packet memory	30000
Force packet plugin	Disabled
Force stream plugin	Disabled
Statistics collection interval	1 (in minutes)
Application System Cache	
Status	Enabled
Max Number of entries in cache	131072
Cache timeout	3600 (in seconds)
Protocol Bundle	
Download Server	https://services.netscreen.com/cgi-bin/index.cgi
AutoUpdate	Disabled
Slot 1:	
Status	Active
Version	1.30.4-22.005 (build date Jan 17 2014)
Sessions	0
Slot 2	
Status	Free

show services application-identification version

Syntax

```
show services application-identification version
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display the Junos OS application package version.

Required Privilege Level

view

RELATED DOCUMENTATION

[request services application-identification download](#) | [1296](#)

List of Sample Output

[show services application-identification version on page 1499](#)

Sample Output

```
show services application-identification version
```

```
user@host> show services application-identification version
```

```
Application package version: 1608
```

show services captive-portal-content-delivery

Syntax

```
show services captive-portal-content-delivery
<pic pic-name>
<profile profile-name>
<rule rule-name> <term term-name>
<ruleset ruleset-name>
<sset sset-name> <brief> <detail> <summary>
<statistics <interface pic-name>>
```

Release Information

Command introduced in Junos OS Release 10.4.

Description

Display the current operational state of all captive portal interfaces.

Options

brief—(Optional) Display brief service set database information.

detail—(Optional) Display detailed service set database information.

pic—Display the PIC database.

profile—Display the profile database.

rule—Display the rule database.

ruleset—Display the rule set database.

sset—Display the service set database.

statistics—Display captive portal content delivery statistics about a PIC.

summary—(Optional) Display a summary of service set database information.

term—(Optional) Display term information for the rule database.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear services captive-portal-content-delivery statistics](#) | 1273

List of Sample Output

[show services captive-portal-content-delivery on page 1502](#)

[show services captive-portal-content-delivery \(Profile\) on page 1503](#)

[show services captive-portal-content-delivery \(Profile HTTP Redirect\) on page 1503](#)

[show services captive-portal-content-delivery \(Profile IPDA Rewrite\) on page 1503](#)

[show services captive-portal-content-delivery \(Rules\) on page 1503](#)

[show services captive-portal-content-delivery \(Rewrite Term\) on page 1503](#)

[show services captive-portal-content-delivery \(Redirect Term\) on page 1504](#)

[show services captive-portal-content-delivery \(Service Set Detail\) on page 1504](#)

[show services captive-portal-content-delivery \(Interface\) on page 1504](#)

[show services captive-portal-content-delivery \(Subscriber with Insert Action Detail\) on page 1504](#)

[show services captive-portal-content-delivery \(Subscriber with Redirect Action Detail\) on page 1505](#)

Output Fields

Table 87 on page 1501 lists the output fields for the **show services captive-portal-content-delivery** command. Output fields are listed in the approximate order in which they appear.

Table 87: show services captive-portal-content-delivery Output Fields

Field Name	Field Description	Level of Output
Name	Name of the interface.	none
Index		none
Profile	Name of the service profile for the HTTP redirect services that contains the rules or rule sets specifying the service.	none
Rules or Rule Sets	List of rules or rule sets contained in the HTTP redirect service profile.	none
Rule Name	Name of an HTTP redirect service rule.	none
Term Name	Name of a rule term.	none
Rule match direction	Traffic direction on the interface where the rule match is applied, input , output , or input-output .	none
Term action	Action performed on packets when the rule term is matched: <ul style="list-style-type: none"> • Accept the packets. • Redirect the packets to a new destination URL. • Rewrite the packets with a new destination address and optionally a new destination port. • Log information about the packet to a system log file. 	none

Table 87: show services captive-portal-content-delivery Output Fields (*continued*)

Field Name	Field Description	Level of Output
Term action option	Additional information related to the term action. <ul style="list-style-type: none"> • A new destination URL for a redirect action. • A new destination address for a rewrite action. • A new destination port for a rewrite action. 	none
Service Sets	Name of service sets contained in a profile.	none
Id	Identifier number for a service set.	none
Compiled Rules		none
service-set interface	Interface on which the service set rules are applied.	none
Packets received	Number of packets received on the service-set interface.	none
Packets altered	Number of packets redirected or rewritten on the service-set interface.	none
Packets dropped	Number of packets dropped on the interface.	detail
Received	Number of packets received for the listed action: Redirect , Rewrite , or Insert .	detail
Altered	Number of packets altered by the listed action: Redirect or Rewrite .	detail
Redirected	Number of packets redirected by the Insert action.	detail
Tag-inserted	Number of packets for which one or more tags has been inserted by the Insert action.	detail

Sample Output

```
show services captive-portal-content-delivery
```

```
user@host> show services captive-portal-content-delivery pic si-5/0/0
```

Name	Index
si-5/0/0	20

show services captive-portal-content-delivery (Profile)

```
user@host> show services captive-portal-content-delivery profile
```

Profile	Rules or Rule Sets
http-redirect	1
ipda-rewrite	1

show services captive-portal-content-delivery (Profile HTTP Redirect)

```
user@host> show services captive-portal-content-delivery profile http-redirect
```

Profile	Rules or Rule Sets
http-redirect	1

show services captive-portal-content-delivery (Profile IPDA Rewrite)

```
user@host> show services captive-portal-content-delivery profile ipda-rewrite
```

Profile	Rules or Rule Sets
ipda-rewrite	1

show services captive-portal-content-delivery (Rules)

```
user@host> show services captive-portal-content-delivery rule
```

Rule Name	Term Name
redirect	t2
rewrite	t1

show services captive-portal-content-delivery (Rewrite Term)

```
user@host> show services captive-portal-content-delivery rule rewrite term t1
```

```
Rule name: rewrite
Rule match direction: input
Term name: t1
Term action: rewrite
```

```
Term action option: null
```

show services captive-portal-content-delivery (Redirect Term)

```
user@host> show services captive-portal-content-delivery rule redirect term t2
```

```
Rule name: redirect
Rule match direction: input
Term name: t2
Term action: redirect
Term action option: http://www.example.net
```

show services captive-portal-content-delivery (Service Set Detail)

```
user@host> show services captive-portal-content-delivery sset sset1 detail
```

Service Set	Id	Profile	Compiled Rules
sset1	1	ipda-rewrite	1

show services captive-portal-content-delivery (Interface)

```
user@host> show services captive-portal-content-delivery statistics interface sis-5/0/0
```

```
service-set interface: si-5/0/0

Packets received  Packets altered
5                3
```

show services captive-portal-content-delivery (Subscriber with Insert Action Detail)

```
user@host> show services captive-portal-content-delivery statistics interface si-5/0/0 detail
```

```
service-set interface: si-5/0/0

Packets dropped: 10

Redirect:
```

Received	Altered	
0	0	
Rewrite:		
Received	Altered	
0	0	
Insert:		
Received	Redirected	Tag-inserted
2	1	1

show services captive-portal-content-delivery (Subscriber with Redirect Action Detail)

user@host> show services captive-portal-content-delivery statistics interface si-4/0/0 detail

service-set interface: si-4/0/0		
Packets dropped: 10		
Redirect:		
Received	Altered	
4	4	
Rewrite:		
Received	Altered	
0	0	
Insert:		
Received	Redirected	Tag-inserted
0	0	0

show services lrf collector statistics

Syntax

```
show services lrf collector statistics
<collector-name>
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display LRF statistics for one or more collectors. If a collector is not specified, statistics are displayed for all collectors.

Options

none—Display LRF statistics for all collectors.

collector-name—(Optional) Display LRF statistics for the specified collector.

Required Privilege Level

view

RELATED DOCUMENTATION

| [Logging and Reporting Function for Subscribers](#) | [441](#)

List of Sample Output

[show services lrf collector statistics on page 1507](#)

Output Fields

[Table 88 on page 1506](#) lists the output fields for the **show services lrf collector statistics** command. Output fields are listed in the approximate order in which they appear.

Table 88: show services lrf collector statistics Output Fields

Field Name	Field Description
Interface	Name of the interface from which data records are sent to the collector.
Templates registered	Number of templates registered with the collector.
Template registration failures	Number of template registration failures.
Templates active	Number of active templates.

Table 88: show services lrf collector statistics Output Fields (*continued*)

Field Name	Field Description
Sessions received	Number of data sessions received for logging of data.
Sessions ignored	Number of data sessions received for logging of data that were ignored.
Records logged	Number of logs sent to the collector.
Records exported	Number of data records exported to the collector.
Record export failures	Number of data record export attempts that failed.

Sample Output

show services lrf collector statistics

user@host> **show services lrf collector statistics**

```
LRF Collector Statistics
  Interface: ms-2/1/0
  Templates registered: 0, Template registration failures: 0, Templates active:
1
  Sessions received: 0, Sessions ignored: 0, Records logged: 0
  Records exported: 0, Record export failures: 0
```

show services lrf rule statistics

Syntax

```
show services lrf rule statistics
<rule-name>
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display LRF statistics for one or more LRF rules. If an LRF rule is not specified, statistics are displayed for all LRF rules.

Options

none—Display LRF statistics for all LRF rules.

rule-name—(Optional) Display LRF statistics for the specified LRF rule.

Required Privilege Level

view

RELATED DOCUMENTATION

| [Logging and Reporting Function for Subscribers](#) | 441

List of Sample Output

[show services lrf rule statistics on page 1509](#)

Output Fields

[Table 89 on page 1508](#) lists the output fields for the **show services lrf rule statistics** command. Output fields are listed in the approximate order in which they appear.

Table 89: show services lrf rule statistics Output Fields

Field Name	Field Description
Interface	Name of the interface from which data records are sent to the collector.
Rule	Name of the LRF rule that caused data records to be exported to the collector.
Template	Name of the template that was used to export data records to the collector.
Templates registered	Number of templates registered with the collector.

Table 89: show services lrf rule statistics Output Fields (*continued*)

Field Name	Field Description
Template registration failures	Number of template registration failures.
Collector	Name of the collector to which data records were sent.
Sessions received	Number of data sessions received for logging of data.
Sessions ignored	Number of data sessions received for logging of data that were ignored.
Sessions logged	Number of data sessions that had data records exported to the collector.
Records exported	Number of data records exported to the collector.
Record export failures	Number of data record export attempts that failed.

Sample Output

show services lrf rule statistics

user@host> **show services lrf rule statistics**

```

LRF Rule Statistics
  Interface: ms-3/1/0
  Rule: r1
  Template: templ
  Templates registered: 2, Template registration failures: 0
  Collector: coll1
  Sessions received: 115, Sessions ignored: 0, Sessions logged: 134
  Records exported: 134, Record export failures: 0

```

show services lrf statistics

Syntax

```
show services lrf statistics
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display number of bytes, packets, and flows for carrying data records to the collector.

Required Privilege Level

view

RELATED DOCUMENTATION

[Logging and Reporting Function for Subscribers](#) | 441

List of Sample Output

[show services lrf statistics on page 1511](#)

Output Fields

[Table 90 on page 1510](#) lists the output fields for the **show services lrf statistics** command. Output fields are listed in the approximate order in which they appear.

Table 90: show services lrf statistics Output Fields

Field Name	Field Description
Interface	Name of the interface from which data records are sent to the collector.
Flow packets	Number of packets carrying data records to the collector.
Flow bytes	Number of bytes carrying data records to the collector.
Active flows	Number of active flows carrying data records to the collector.
Total flows	Total number of flows for carrying data records to the collector.

Sample Output

show services lrf statistics

user@host> **show services lrf statistics**

```
LRF Statistics
  Interface: ms-3/1/0
  Flow packets: 31125, Flow bytes: 15335751
  Active flows: 0, Total flows: 1887

  Interface: ms-3/2/0
  Flow packets: 0, Flow bytes: 0
  Active flows: 0, Total flows: 0
```

show services lrf template

Syntax

```
show services lrf template option
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

Display the fields for a template type. You must specify a template type.

Options

option—Specify one of the following template types:

- device-data—Display the fields for the Device Data template type.
- flow-id—Display the fields for the Flow ID template type.
- http—Display the fields for the HTTP template type.
- ifl-subscriber—Display the fields for the IFL Subscriber template type.
- ipflow—Display the fields for the IPFlow template type.
- ipflow-extended—Display the fields for the IPFlow Extended template type.
- ipflow-tcp—Displays the fields for the IPFlow TCP template type.
- ipflow-tcp-ts—Displays the fields for the IPFlow TCP Timestamp template type.
- ipflow-ts—Display the fields for the IPFlow Timestamp template type.
- ipv4—Display the fields for the IPv4 template type.
- ipv4-extended—Display the fields for the IPv4 Extended template type.
- ipv6—Display the fields for the IPv6 template type.
- ipv6-extended—Display the fields for the IPv6 Extended template type.
- l7-app—Display the fields for the L7 Application template type.
- mobile-subscriber—Display the fields for the Mobile Subscriber template type.
- pcc—Display the fields for the PCC template type.
- subscriber-data—Display the fields for the Subscriber Data template type.
- wireline-subscriber—Display the fields for the Wireline Subscriber template type.

Required Privilege Level

view

RELATED DOCUMENTATION

| [Logging and Reporting Function for Subscribers](#) | 441

List of Sample Output

- [show services lrf template ipv4 on page 1513](#)
- [show services lrf template ipflow-extended on page 1513](#)
- [show services lrf template ipflow-tcp-ts on page 1513](#)
- [show services lrf template ipflow-tcp on page 1514](#)

Sample Output

show services lrf template ipv4

user@host> show services lrf template ipv4

```
LRF Template fields
  Ipv4 source address
  Ipv4 destination address
  TCP/UDP source port
  TCP/UDP destination port
```

show services lrf template ipflow-extended

user@host> show services lrf template ipflow-extended

Field	Element Id	Length(bytes)	Vendor
Service set name	520	16	Juniper
Routing-instance	521	16	Juniper

show services lrf template ipflow-tcp-ts

user@host> show services lrf template ipflow-tcp-ts

Field	Element Id	Length(bytes)	Vendor
Smooth RTT uplink	10000	4	Juniper
Smooth RTT downlink	10001	4	Juniper
Client setup Time	10002	4	Juniper
Server Setup time	10003	4	Juniper
Client first payload timestamp	10004	8	Juniper
Upload time	10005	4	Juniper
Server first payload timestamp	10006	8	Juniper

Download time	10007	4	Juniper
Acknowledged volumes uplink	10008	8	Juniper
Acknowledged volumes downlink	10009	8	Juniper

show services lrf template ipflow-tcp

```
user@host> show services lrf template ipflow-tcp
```

Field	Element Id	Length(bytes)	Vendor
Retransmitted TCP packets uplink	115	4	Juniper
Retransmitted TCP packets downlink	116	4	Juniper
TCP flow creation timestamp	121	8	Juniper

show services pcef pic

Syntax

```
show services pcef pic
<fpc-slot slot-number>
<pic-slot pic-number>
```

Release Information

Command introduced in Junos OS Release 17.2 on MX Series routers.

Description

Displays the number of Junos OS Subscriber Management subscribers present on each service PIC who are a using policy control and enforcement function (PCEF) profile to define the treatment to apply to packets associated with specific applications (for example, Facebook) or to specific service data flows.

Options

fpc-slot slot-number—(Optional) Display the number of PCEF subscribers present on the specified FPC.

pic-slot pic-number—(Optional) Display the number of PCEF subscribers present on the specified PIC.

Required Privilege Level

view

RELATED DOCUMENTATION

| [show services pcef subscribers](#) | [1517](#)

List of Sample Output

[show services pcef pic on page 1516](#)

Output Fields

[Table 91 on page 1515](#) lists the output fields for the **show services pcef pic** command. Output fields are listed in the approximate order in which they appear.

Table 91: show services pcef pic Output Fields

Field Name	Field Description
FPC Slot	FPC slot on which the PCEF subscribers are present.
PIC Slot	PIC slot on which the PCEF subscribers are present.
Active Subscribers	Number of active PCEF subscribers that are present on the PIC slot.

Sample Output

show services pcef pic

user@host> **show services pcef pic**

FPC Slot	PIC Slot	Active Subscribers
2	0	1
2	1	1
2	2	0
2	3	0

show services pcef subscribers

Syntax

```
show services pcef subscribers
<detail | extensive | summary | terse>
<aci-interface-set-name aci-interface-set-name>
<address address>
<agent-circuit-identifier agent-circuit-identifier-substring>
<client-type client-type>
<count>
<id>
<interface interface>
<logical-system logical-system>
<mac-address mac-address>
<physical-interface physical-interface-name>
<profile-name profile-name>
<routing-instance routing-instance>
<stacked-vlan-id stacked-vlan-id>
<subscriber-state subscriber-state>
<user-name user-name>
<vci vci-identifier>
<vpi vpi-identifier>
<vlan-id vlan-id>
```

Release Information

Command introduced in Junos OS Release 17.2 on MX Series routers.

Description

Displays information for Junos OS Subscriber Management subscribers who are using a policy control and enforcement function (PCEF) profile to define the treatment to apply to packets associated with specific applications (for example, Facebook) or to specific service data flows.

Options

detail | extensive | summary | terse—(Optional) Display the specified level of output.

aci-interface-set-name—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

address—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier-substring—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

client-type—(Optional) Display subscribers whose client type matches one of the following client types:

- **dhcp**—DHCP clients only.
- **dot1x**—Dot1x clients only.
- **essm**—ESSM clients only.
- **fwauth**—FwAuth (authenticated across a firewall) clients only.
- **l2tp**—L2TP clients only.
- **mlppp**—MLPPP clients only.
- **ppp**—PPP clients only.
- **pppoe**—PPPoE clients only.
- **static**—Static clients only.
- **vlan**—VLAN clients only.
- **vpls-pw**—VPLS pseudowire clients only.
- **xauth**—Xauth clients only.

count—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the **count** option alone or with the **address**, **client-type**, **interface**, **logical-system**, **mac-address**, **profile-name**, **routing-instance**, **stacked-vlan-id**, **subscriber-state**, or **vlan-id** options.

id—(Optional) Display a specific subscriber session whose session id matches the specified subscriber ID. You can display subscriber IDs by using the **show subscribers extensive** or the **show subscribers interface extensive** commands.

interface—(Optional) Display subscribers whose interface matches the specified interface.

logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.

mac-address—(Optional) Display subscribers whose MAC address matches the specified MAC address.

physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

profile-name—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

subscriber-state—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

vci-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is 0 through 255.

vpi-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is 0 through 65,535.

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the **stacked-vlan-id stacked-vlan-id** option to match the outer VLAN tag.

NOTE: Because of display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

RELATED DOCUMENTATION

| [show services pcef pic](#) | [1515](#)

List of Sample Output

[show services pcef subscribers terse on page 1522](#)

[show services pcef subscribers extensive on page 1523](#)

[show services pcef subscribers summary on page 1523](#)

Output Fields

[Table 92 on page 1520](#) lists the output fields for the **show services pcef subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 92: show services pcef subscribers Output Fields

Field Name	Field Description	Level of Output
Interface	Interface associated with the subscriber.	detail extensive terse
IP Address	Subscriber IPv4 or IPv6 address.	detail extensive terse
User Name	Name of subscriber.	detail extensive terse
Packets	Number of the subscriber's packets that are processed by a PCEF profile. <ul style="list-style-type: none"> • Input—Input packets. • Output—Output packets. 	terse
Packet Drops	Number of the subscriber's packets that were dropped as a result of being processed by a PCEF profile. <ul style="list-style-type: none"> • Input—Input packets. • Output—Output packets. 	terse
Session ID	ID number for a subscriber service session.	extensive detail
PFE Flow ID	Variable-based forwarding flow ID.	extensive detail
PCEF profile	PCEF profile that is assigned to the subscriber.	extensive detail
Routing Instance	Routing instance associated with the subscriber.	extensive detail

Table 92: show services pcef subscribers Output Fields (*continued*)

Field Name	Field Description	Level of Output
Service VRF		extensive detail
Service set	Service set that is performing policy control.	extensive detail
Input packets	Number of the subscriber's input packets that are processed by a PCEF profile.	extensive detail
Input octet	Number of the subscriber's input octets that are processed by a PCEF profile.	extensive detail
Output packets	Number of the subscriber's output packets that are processed by a PCEF profile.	extensive detail
Output octet	Number of the subscriber's output octets that are processed by a PCEF profile.	extensive detail
Input drops	Number of the subscriber's input packets that were dropped as a result of being processed by a PCEF profile.	extensive detail
Input drop bytes	Number of the subscriber's input bytes that were dropped as a result of being processed by a PCEF profile.	extensive detail
Output drops	Number of the subscriber's output packets that were dropped as a result of being processed by a PCEF profile.	extensive detail
Output drop bytes	Number of the subscriber's output bytes that were dropped as a result of being processed by a PCEF profile.	extensive detail
Rule count	Number of PCC rules that were applied to subscriber's traffic.	extensive detail

Table 92: show services pcef subscribers Output Fields (*continued*)

Field Name	Field Description	Level of Output
Rule information	<p>For each PCC rule that is assigned to the subscriber, the following information appears:</p> <ul style="list-style-type: none"> • Rule name—Name of PCC rule. • In sessions—Number of incoming sessions for the subscriber that are processed by the rule. • Out sessions—Number of outgoing sessions for the subscriber that are processed by the rule. • Input packets—Number of input packets and bytes for the subscriber that are processed by the rule. • Output packets—Number of output packets and bytes for the subscriber that are processed by the rule. • Input drop packets—Number of input packets and bytes for the subscriber that are dropped by the rule. • Output drop packets—Number of output packets and bytes for the subscriber that are dropped by the rule. 	extensive detail
Total subscribers	Total number of subscribers.	summary

Sample Output

```
show services pcef subscribers terse
```

```
user@host> show services pcef subscribers terse
```

Interface Name	IP Address	User Name	Packets		Packet Drops	
			Input	Output	Input	Output
demux0.3221225518	192.0.2.26	pcefuser	4215521	5155789	0	589900

show services pcef subscribers extensive

```
user@host> show services pcef subscribers extensive
```

```

Session ID: 60
PFE flow ID: 73
Interface: demux0.3221225518
IP address: 192.0.2.26
Username: pcefuser
PCEF profile: pcef-prof-1
Routing Instance: default
Service VRF: 1
Service set: 0
Input packets: 4229161
Input octet: 1632292118
Output packets: 5171863
Output octet: 2199274163
Input drops: 0
Input drop bytes: 0
Output drops: 591751
Output drop bytes: 882496606
Rule count: 2
  Rule information      :
    Rule name          :          limit-fb
    In sessions         :             13680
    Out sessions        :              0
    Input packets       :             2629243      244119982 bytes
    Output packets      :             3702363      1117717893 bytes
    Input drop packets  :              0           0 bytes
    Output drop packets :             591751      882496606 bytes
    Rule name           :          default
    In sessions         :             84545
    Out sessions        :              0
    Input packets       :             1599918      1388172136 bytes
    Output packets      :             1469500      1081556270 bytes
    Input drop packets  :              0           0 bytes
    Output drop packets :              0           0 bytes

```

Sample Output

show services pcef subscribers summary

```
user@host> show services pcef subscribers summary
```

Total subscribers: 1

show services service-sets summary

Syntax

```
show services service-sets summary
<interface interface-name>
```

Release Information

Command introduced before Junos OS Release 7.4.
 Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display service set summary information.

Options

- none**—Display service set summary information for all adaptive services interfaces.
- interface *interface-name***—(Optional) Display service set summary information for a particular interface.
 On M Series and T Series routers, *interface-name* can be *ms-fpc/pic/port*, *sp-fpc/pic/port*, or *rspnumber*.

 On MX Series MX240, MX480, and MX960 routers, *interface-name* can be *vms-fpc/pic/port* for the MX-SPC3 services card for Next Gen Services.

Required Privilege Level

view

List of Sample Output

- [show services service-sets summary on page 1526](#)
- [show services service-sets summary interface on page 1526](#)

Output Fields

[Table 93 on page 1525](#) lists the output fields for the **show services service-sets summary** command. Output fields are listed in the approximate order in which they appear.

Table 93: show services service-sets summary Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface
Service type	Type of adaptive service, such as stateful firewall (SFW), Network Address Translation (NAT), intrusion detection service (IDS), Layer 2 Tunneling Protocol (L2TP), Compressed Real-Time Transport Protocol (CRTP), or IP Security (IPsec)

Table 93: show services service-sets summary Output Fields (*continued*)

Field Name	Field Description
Service sets configured	Total number of service sets configured on the PIC that use internal service set IDs and do not consume external service sets, including CRTP and L2TP
Bytes used	Bytes used by a particular service or all services
Policy bytes used	Policy bytes used by a particular service or all services
CPU utilization	Percentage of the CPU resources being used

Sample Output

show services service-sets summary

user@host> **show services service-sets summary**

Service sets				
Interface	CPU configured	Bytes used	Session bytes used	Policy
bytes used	utilization			
vms-3/0/0	1	3453621040 (24.93%)	0 (0.00%)	8161168
(0.90%)	0.14 %			

show services service-sets summary interface

user@host> **show services service-sets summary interface sp-1/3/0**

Interface: sp-1/3/0				
Service sets				CPU
Service type	configured	Bytes used		utilization
SFW/NAT/IDS	1	54 (0.00 %)		N/A
L2TP	1	58 (0.00 %)		N/A
CRTP	1	58 (0.00 %)		N/A
System	0	920831 (0.44 %)		N/A
Idle	0	0 (0.00 %)		N/A
Total	3	921001 (0.44 %)		N/A

show subscribers

Syntax

```
show subscribers
<detail | extensive | terse>
<aci-interface-set-name aci-interface-set-name>
<address address>
<agent-circuit-identifier agent-circuit-identifier>
<agent-remote-identifier agent-remote-identifier>
<aggregation-interface-set-name interface-set-name>
<client-type client-type>
<count>
<id session-id <accounting-statistics>>
<interface interface <accounting-statistics>>
<logical-system logical-system>
<mac-address mac-address>
<physical-interface physical-interface-name>
<profile-name profile-name>
<routing-instance routing-instance>
<stacked-vlan-id stacked-vlan-id>
<subscriber-state subscriber-state>
<user-name user-name>
<vci vci-identifier>
<vpi vpi-identifier>
<vlan-id vlan-id>
```

Release Information

Command introduced in Junos OS Release 9.3.

Command introduced in Junos OS Release 9.3 for EX Series switches.

client-type, **mac-address**, **subscriber-state**, and **extensive** options introduced in Junos OS Release 10.2.

count option usage with other options introduced in Junos OS Release 10.2.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Options **aci-interface-set-name** and **agent-circuit-identifier** introduced in Junos OS Release 12.2.

The **physical-interface** and **user-name** options introduced in Junos OS Release 12.3.

Options **vci** and **vpi** introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.

Options **vci** and **vpi** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Enhanced subscriber management supported in Junos OS Release 15.1R3 on MX Series routers.

accounting-statistics option added in Junos OS Release 15.1R3 and 17.4R1 on MX Series routers.

aggregation-interface-set-name option added in Junos OS Release 18.4R1 on MX Series routers.

Description

Display information for active subscribers.

Options

detail | extensive | terse—(Optional) Display the specified level of output.

aci-interface-set-name—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

address—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified string. You can specify either the complete ACI string or a substring. To specify a substring, you must enter characters that form the beginning of the string, followed by an asterisk (*) as a wildcard to substitute for the remainder of the string. The wildcard can be used only at the end of the specified substring; for example:

```
user@host1> show subscribers agent-circuit-identifier substring*
```

Junos OS Release	Substring Support
Junos OS Release 13.3R1	You can specify a substring without a wildcard.
Starting in Junos OS Release 14.1R1	You must specify the complete ACI string; you cannot specify a wildcard.
Starting in Junos OS Release 15.1R7, 16.1R7, 16.2R3, 17.1R3, 17.2R3, 17.3R3, 17.4R2, 18.1R2, 18.2R1	You can specify a substring, but you must include the wildcard character at the end of the substring.

agent-remote-identifier—(Optional) Display all dynamic subscriber sessions whose ARI value matches the specified string. You must specify the complete ACI string; you cannot specify a wildcard.

aggregation-interface-set-name interface-set-name—(Optional) Display summary information for the specified aggregation node interface set, including interface, VLAN ID, username and LS:RI.

client-type—(Optional) Display subscribers whose client type matches one of the following client types:

- **dhcp**—DHCP clients only.
- **dotlx**—Dotlx clients only.
- **essm**—ESSM clients only.

- **fixed-wireless-access**—Fixed wireless access clients only.
- **fwauth**—FwAuth (authenticated across a firewall) clients only.
- **hag-bundle**—HAG tunnel bundle clients only.
- **hag-tunnel**—HAG tunnel clients only.
- **l2tp**—L2TP clients only.
- **mlppp**—MLPPP clients only.
- **ppp**—PPP clients only.
- **pppoe**—PPPoE clients only.
- **static**—Static clients only.
- **vlan**—VLAN clients only.
- **vlan-oob**—VLAN out-of-band (ANCP-triggered) clients only.
- **vpls-pw**—VPLS pseudowire clients only.
- **xauth**—Xauth clients only.

count—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the **count** option alone or with the **address**, **client-type**, **interface**, **logical-system**, **mac-address**, **profile-name**, **routing-instance**, **stacked-vlan-id**, **subscriber-state**, or **vlan-id** options.

id session-id—(Optional) Display a specific subscriber session whose session ID matches the specified subscriber ID. You can display subscriber IDs by using the **show subscribers extensive** or the **show subscribers interface extensive** commands.

id session-id accounting-statistics—(Optional) Display accurate subscriber accounting statistics for a subscriber session with the specified ID. Requires the **actual-transmit-statistics** statement to be configured in the dynamic profile for the dynamic logical interface. If the statement is not configured, a value of 0 is displayed for accounting statistics.

interface—(Optional) Display subscribers whose interface matches the specified interface.

interface accounting-statistics—(Optional) Display subscriber accounting statistics for the specified interface. Requires the **actual-transmit-statistics** statement to be configured in the dynamic profile for the dynamic logical interface.

logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.

mac-address—(Optional) Display subscribers whose MAC address matches the specified MAC address.

physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

profile-name—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

subscriber-state—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

vci-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is 0 through 255.

vpi-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is 0 through 65,535.

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the **stacked-vlan-id** **stacked-vlan-id** option to match the outer VLAN tag.

NOTE: Because of display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

RELATED DOCUMENTATION

[show subscribers summary](#) | 1576

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers

Verifying and Managing Junos OS Enhanced Subscriber Management

List of Sample Output

[show subscribers \(IPv4\) on page 1542](#)

[show subscribers \(IPv6\) on page 1542](#)
[show subscribers \(IPv4 and IPv6 Dual Stack\) on page 1542](#)
[show subscribers \(Single Session DHCP Dual Stack\) on page 1543](#)
[show subscribers \(Single Session DHCP Dual Stack detail\) on page 1543](#)
[show subscribers \(LNS on MX Series Routers\) on page 1544](#)
[show subscribers \(L2TP Switched Tunnels\) on page 1544](#)
[show subscribers aggregation-interface-set-name on page 1544](#)
[show subscribers client-type dhcp detail on page 1544](#)
[show subscribers client-type dhcp detail \(DHCPv6\) on page 1545](#)
[show subscribers client-type dhcp extensive on page 1546](#)
[show subscribers client-type fixed-wireless-access on page 1547](#)
[show subscribers client-type fixed-wireless-access detail \(Detail\) on page 1547](#)
[show subscribers client-type hag-bundle on page 1548](#)
[show subscribers client-type hag-bundle \(Detail\) on page 1548](#)
[show subscribers client-type -hag-tunnel on page 1548](#)
[show subscribers client-type hag-tunnel \(Detail\) on page 1549](#)
[show subscribers client-type vlan-oob detail on page 1549](#)
[show subscribers count on page 1550](#)
[show subscribers address detail \(IPv6\) on page 1550](#)
[show subscribers detail \(IPv4\) on page 1551](#)
[show subscribers detail \(IPv6\) on page 1551](#)
[show subscribers detail \(pseudowire Interface for GRE Tunnel\) on page 1552](#)
[show subscribers detail \(IPv6 Static Demux Interface\) on page 1552](#)
[show subscribers detail \(L2TP LNS Subscribers on MX Series Routers\) on page 1552](#)
[show subscribers detail \(L2TP Switched Tunnels\) on page 1553](#)
[show subscribers detail \(Tunneled Subscriber\) on page 1554](#)
[show subscribers detail \(IPv4 and IPv6 Dual Stack\) on page 1554](#)
[show subscribers detail \(ACI Interface Set Session\) on page 1555](#)
[show subscribers detail \(PPPoE Subscriber Session with ACI Interface Set\) on page 1556](#)
[show subscribers extensive on page 1556](#)
[show subscribers extensive \(Aggregation Node Interface Set and DSL Forum Attributes\) on page 1557](#)
[show subscribers extensive \(Passive Optical Network Circuit Interface Set\) on page 1558](#)
[show subscribers extensive \(DNS Addresses from Access Profile or Global Configuration\) on page 1559](#)
[show subscribers extensive \(DNS Addresses from RADIUS\) on page 1560](#)
[show subscribers extensive \(IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration\) on page 1560](#)
[show subscribers extensive \(RPF Check Fail Filter\) on page 1561](#)
[show subscribers extensive \(L2TP LNS Subscribers on MX Series Routers\) on page 1561](#)
[show subscribers extensive \(IPv4 and IPv6 Dual Stack\) on page 1562](#)
[show subscribers extensive \(ADF Rules \) on page 1563](#)
[show subscribers extensive \(Effective Shaping-Rate\) on page 1564](#)
[show subscribers extensive \(PPPoE Subscriber Access Line Rates on page 1564](#)
[show subscribers extensive \(Subscriber Session Using PCEF Profile\) on page 1566](#)

[show subscribers aci-interface-set-name detail \(Subscriber Sessions Using Specified ACI Interface Set\) on page 1567](#)

[show subscribers agent-circuit-identifier detail \(Subscriber Sessions Using Specified ACI Substring\) on page 1568](#)

[show subscribers id accounting-statistics on page 1569](#)

[show subscribers interface accounting-statistics on page 1569](#)

[show subscribers interface extensive on page 1570](#)

[show subscribers logical-system terse on page 1571](#)

[show subscribers physical-interface count on page 1571](#)

[show subscribers routing-instance inst1 count on page 1572](#)

[show subscribers stacked-vlan-id detail on page 1572](#)

[show subscribers stacked-vlan-id vlan-id detail \(Combined Output\) on page 1572](#)

[show subscribers stacked-vlan-id vlan-id interface detail \(Combined Output for a Specific Interface\) on page 1572](#)

[show subscribers user-name detail on page 1573](#)

[show subscribers vlan-id on page 1573](#)

[show subscribers vlan-id detail on page 1573](#)

[show subscribers vpi vci extensive \(PPPoE-over-ATM Subscriber Session\) on page 1574](#)

[show subscribers address detail \(Enhanced Subscriber Management\) on page 1574](#)

Output Fields

[Table 94 on page 1532](#) lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 94: show subscribers Output Fields

Field Name	Field Description
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session.
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions the value is Tunnel-switched .
User Name	Name of subscriber.
LS:RI	Logical system and routing instance associated with the subscriber.
Type	Subscriber client type (DHCP, FWA, GRE, HAG-BUNDLE, HAG-TUNNEL, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).

Table 94: show subscribers Output Fields (*continued*)

Field Name	Field Description
IP Address	Subscriber IPv4 address.
IP Netmask	Subscriber IP netmask. (MX Series) This field displays 255.255.255.255 by default. For tunneled or terminated PPP subscribers only, this field displays the actual value of Framed-IP-Netmask when the SDB_FRAMED_PROTOCOL attribute in the session database is equal to AUTHD_FRAMED_PROTOCOL_PPP. This occurs in the use case where the LNS generates access-internal routes when it receives Framed-IP-Netmask from RADIUS during authorization. When it receives Framed-Pool from RADIUS, the pool mask is ignored and the default /32 mask is used.
Primary DNS Address	IP address of primary DNS server. This field is displayed with the extensive option only when the address is provided by RADIUS.
Secondary DNS Address	IP address of secondary DNS server. This field is displayed with the extensive option only when the address is provided by RADIUS.
IPv6 Primary DNS Address	IPv6 address of primary DNS server. This field is displayed with the extensive option only when the address is provided by RADIUS.
IPv6 Secondary DNS Address	IPv6 address of secondary DNS server. This field is displayed with the extensive option only when the address is provided by RADIUS.
Domain name server inet	IP addresses for the DNS server, displayed in order of configuration. This field is displayed with the extensive option only when the addresses are derived from the access profile or the global access configuration.
Domain name server inet6	IPv6 addresses for the DNS server, displayed in order of configuration. This field is displayed with the extensive option only when the addresses are derived from the access profile or the global access configuration.
Primary WINS Address	IP address of primary WINS server.
Secondary WINS Address	IP address of secondary WINS server.

Table 94: show subscribers Output Fields (*continued*)

Field Name	Field Description
IPv6 Address	Subscriber IPv6 address, or multiple addresses.
IPv6 Prefix	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
IPv6 User Prefix	IPv6 prefix obtained through NDRA.
IPv6 Address Pool	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
IPv6 Network Prefix Length	Length of the network portion of the IPv6 address.
IPv6 Prefix Length	Length of the subscriber IPv6 prefix.
Logical System	Logical system associated with the subscriber.
Routing Instance	Routing instance associated with the subscriber.
Interface	(Enhanced subscriber management for MX Series routers) Name of the enhanced subscriber management logical interface, in the form demux0.nnnn (for example, demux0.3221225472), to which access-internal and framed subscriber routes are mapped.
Interface Type	Whether the subscriber interface is Static or Dynamic .

Table 94: show subscribers Output Fields (*continued*)

Field Name	Field Description
Interface Set	<p>Internally generated name of the dynamic ACI or ALI interface set used by the subscriber session. The prefix of the name indicates the string received in DHCP or PPPoE control packets on which the interface set is based. For ALI interface sets, the prefix indicates that the value is configured as a trusted option to identify the subscriber line.</p> <p>The name of the interface set uses one of the following prefixes:</p> <ul style="list-style-type: none"> • aci—ACI; for example, aci-1033-demux0.3221225524. This is the only prefix allowed for ACI interface sets. • ari—ARI; for example, ari-1033-demux0.3221225524. • aci+ari—Both the ACI and ARI; for example, aci+ari-1033-demux0.3221225524. • noids—Neither the ACI nor the ARI were received; for example, noids-1033-demux0.3221225524. <p>NOTE: ACI interface sets are configured with the agent-circuit-identifier autoconfiguration stanza. ALI interface sets are configured with the line-identity autoconfiguration stanza.</p> <p>Besides dynamic ACI and ALI interface sets, this field can be an interface set based on a substring of the ARI string. This occurs when the dynamic profile includes the predefined variable \$junos-pon-id-interface-set-name, and the profile is applied for a passive optical network (PON). The ARI string is inserted by the optical line terminal (OLT). The final substring in the string, unique for the PON, identifies individual subscriber circuits, and is used as the name of the interface set.</p>
Interface Set Type	Interface type of the ACI interface set: Dynamic . This is the only ACI interface set type currently supported.
Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.
Underlying Interface	Name of the underlying interface for the subscriber session.
Dynamic Profile Name	<p>Dynamic profile used for the subscriber.</p> <p>For a HAG bundle, the profile that creates the bundle interface, which represents the subscriber interface. For a HAG tunnel, the profile the creates the DSL or LTE tunnel interface.</p>
Dynamic Profile Version	Version number of the dynamic profile used for the subscriber.
MAC Address	MAC address associated with the subscriber.
State	Current state of the subscriber session (Init , Configured , Active , Terminating , Tunneled).

Table 94: show subscribers Output Fields (*continued*)

Field Name	Field Description
L2TP State	Current state of the L2TP session, Tunneled or Tunnel-switched . When the value is Tunnel-switched , two entries are displayed for the subscriber; the first entry is at the LNS interface on the LTS and the second entry is at the LAC interface on the LTS.
Tunnel switch Profile Name	Name of the L2TP tunnel switch profile that initiates tunnel switching.
Local IP Address	IP address of the local gateway (LAC).
Remote IP Address	IP address of the remote peer (LNS).
PFE Flow ID	Forwarding flow identifier.
VLAN Id	VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	<p>For the dhcp client type, option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.</p> <p>For the vlan-oob client type, the agent circuit ID or access-loop circuit identifier that identifies the subscriber line based on the subscriber-facing DSLAM interface on which the subscriber request originates.</p>
Agent Remote ID	<p>For the dhcp client type, option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.</p> <p>For the vlan-oob client type, the agent remote ID or access-loop remote identifier that identifies the subscriber line based on the NAS-facing DSLAM interface on which the subscriber request originates.</p>

Table 94: show subscribers Output Fields (continued)

Field Name	Field Description
Aggregation Interface-set Name	<p>Value of the \$junos-aggregation-interface-set-name predefined variable; one of the following:</p> <ul style="list-style-type: none"> When the hierarchical-access-network-detection option is configured for the access lines and the value of the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003) received either in the ANCP Port Up message or PPPoE PADR IA tags begins with a # character, then the variable takes the value of the remainder of the string after the # character. When the hierarchical-access-network-detection option is not configured, or if the sting does not begin with the # character, then the variable takes the value specified with the predefined-variable-defaults statement.
Accounting Statistics	Actual transmitted subscriber accounting statistics by session ID or interface. Service accounting statistics are not included. These statistics do not include overhead bytes or dropped packets; they are the accurate statistics used by RADIUS. The statistics are counted when the actual-transmit-statistics statement is included in the dynamic profile.
DHCP Relay IP Address	IP address used by the DHCP relay agent.
ATM VPI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
ATM VCI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.
Login Time	Date and time at which the subscriber logged in.
DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
Server DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options.
Server DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
DHCPV6 Header	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.

Table 94: show subscribers Output Fields (*continued*)

Field Name	Field Description
IPv4 Input Service Set	Input service set in access dynamic profile.
IPv4 Output Service Set	Output service set in access dynamic profile.
PCEF Profile	PCEF profile in access dynamic profile.
PCEF Rule/Rulebase	PCC rule or rulebase used in dynamic profile.
Dynamic configuration	Values for variables that are passed into the dynamic profile from RADIUS.
Service activation time	Time at which the first family in this service became active.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber session.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session ID	ID number for a subscriber service session.
Service Session Name	Service session profile name.
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.

Table 94: show subscribers Output Fields (*continued*)

Field Name	Field Description
Idle Timeout (seconds)	Number of seconds subscriber can be idle before the session is automatically terminated.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute. This field is displayed only when the predefined variable \$junos-ipv6-address is used in the dynamic profile.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).
IPv6 Input Filter Name	Name assigned to the IPv6 input filter (client or service session).

Table 94: show subscribers Output Fields (*continued*)

Field Name	Field Description
IPv6 Output Filter Name	Name assigned to the IPv6 output filter (client or service session).
IFL Input Filter Name	Name assigned to the logical interface input filter (client or service session).
IFL Output Filter Name	Name assigned to the logical interface output filter (client or service session).
DSL type	PPPoE subscriber's access line type reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute DSL-Type (0x0091). The DSL type is one of the following types: ADSL , ADSL2 , ADSL2+ , OTHER , SDSL , VDSL , or VDSL2 .
Frame/Cell Mode	<p>Mode type of the PPPoE subscriber's access line determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091):</p> <ul style="list-style-type: none"> • Cell—When the DSL line type is one of the following: ADSL, ADSL2, or ADSL2+. • Frame—When the DSL line type is one of the following: OTHER, SDSL, VDSL, or VDSL2. <p>The value is stored in the subscriber session database.</p>
Overhead accounting bytes	Number of bytes added to or subtracted from the actual downstream cell or frame overhead to account for the technology overhead of the DSL line type. The value is determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091). The value is stored in the subscriber session database.
Actual upstream data rate	Unadjusted upstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Upstream (0x0081).
Actual downstream data rate	Unadjusted downstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Downstream (0x0082).
Adjusted downstream data rate	Adjusted downstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.
Adjusted upstream data rate	Adjusted upstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.

Table 94: show subscribers Output Fields (*continued*)

Field Name	Field Description
Local TEID-U	<p>Tunnel endpoint identifier on the BNG for the GTP-U user plane tunnel to the eNodeB. The identifier is allocated by the BNG.</p> <p>A fully qualified local TEID-C consists of this identifier and the GTPU Tunnel Local IP address value.</p>
Local TEID-C	<p>Tunnel endpoint identifier on the BNG for the GTP-C control plane tunnel to the MME. The identifier is allocated by the BNG.</p> <p>A fully qualified local TEID-C consists of this identifier and the GTPC Local IP address value.</p>
Remote TEID-U	<p>Tunnel endpoint identifier on the eNodeB for the GTP-U user plane tunnel to the BNG. The identifier is allocated by the eNodeB.</p> <p>A fully qualified remote TEID-U consists of this identifier and the GTPU Tunnel Remote IP address value.</p>
Remote TEID-C	<p>Tunnel endpoint identifier on the MME for the GTP-C control plane tunnel to the BNG. The identifier is allocated by the MME.</p> <p>A fully qualified remote TEID-C consists of this identifier and the GTPC Remote IP address value.</p>
GTPU Tunnel Remote IP address	<p>IP address of the S1-U interface on the eNodeB for the GTP-U tunnel endpoint.</p> <p>A fully qualified remote TEID-U consists of this address and the Remote TEID-U value.</p>
GTPU Tunnel Local IP address	<p>IP address of the S1-U interface on the BNG for the GTP-U tunnel endpoint.</p> <p>A fully qualified local TEID-U consists of this address and the Local TEID-U value</p>
GTPC Remote IP address	<p>IP address of the S11 interface on the MME for the GTP-C tunnel endpoint.</p> <p>A fully qualified remote TEID-C consists of this address and the Remote TEID-C value.</p>
GTPC Local IP address	<p>IP address of the S11 interface on the BNG for the GTP-C tunnel endpoint.</p> <p>A fully qualified local TEID-C consists of this address and the Local TEID-C value.</p>
Access Point Name	<p>Access point name (APN) for the user equipment. The APN corresponds to the connection and service parameters that the subscriber's mobile device can use for connecting to the carrier's gateway to the Internet.</p>

Table 94: show subscribers Output Fields (continued)

Field Name	Field Description
HAG Session ID	ID number for the hybrid access gateway session. Sharing this ID bonds the DSL and LTE tunnels that make up the bundle.
HAG Tunnel Type	Type of HAG tunnel, DSL or LTE.

Sample Output

show subscribers (IPv4)

```
user@host> show subscribers
```

Interface	IP Address/VLAN ID	User Name	LS:RI
ge-1/3/0.1073741824	10		default:default
demux0.1073741824	203.0.113.10	WHOLESALE-CLIENT	default:default
demux0.1073741825	203.0.113.3	RETAILER1-CLIENT	test1:retailer1
demux0.1073741826	203.0.113.3	RETAILER2-CLIENT	test1:retailer2

show subscribers (IPv6)

```
user@host> show subscribers
```

Interface	IP Address/VLAN ID	User Name	LS:RI
ge-1/0/0.0	2001:db8:c0:0:0:0/74	WHOLESALE-CLIENT	default:default
*	2001:db8:1/128	subscriber-25	default:default

show subscribers (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers
```

Interface	IP Address/VLAN ID	User Name	LS:RI
demux0.1073741834	0x8100.1002 0x8100.1		
default:default			
demux0.1073741835	0x8100.1001 0x8100.1		
default:default			
pp0.1073741836	203.0.113.13	dualstackuser1@example1.com	
default:ASP-1			
*	2001:db8:1::/48		
*	2001:db8:1:1::/64		

```

pp0.1073741837      203.0.113.33      dualstackuser2@example1.com
default:ASP-1
*                   2001:db8:1:2:5::/64

```

show subscribers (Single Session DHCP Dual Stack)

```
user@host> show subscribers
```

Interface	IP Address/VLAN ID	User Name	LS:RI
demux0.1073741364	192.168.10.10	dual-stack-retail35	
default:default	2001:db8::100:0:0:0/74		
default:default	2001:db8:3ffe:0:4::/64		

show subscribers (Single Session DHCP Dual Stack detail)

```
user@host> show subscribers id 27 detail
```

```

Type: DHCP
User Name: dual-stack-retail33
IP Address: 10.10.0.53
IPv6 Address: 2001:db8:3000:0:0:8003::2
IPv6 Prefix: 2001:db8:3ffe:0:4::/64
Logical System: default
Routing Instance: default
Interface: ae0.3221225472
Interface type: Static
Underlying Interface: ae0.3221225472
Dynamic Profile Name: dhcp-retail-18
MAC Address: 00:00:5E:00:53:02
State: Active
DHCP Relay IP Address: 10.10.0.1
Radius Accounting ID: 27
Session ID: 27
PFE Flow ID: 2
Stacked VLAN Id: 2000
VLAN Id: 1
Login Time: 2014-05-15 10:12:10 PDT
DHCP Options: len 60
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 00 64 01 01 02

```

```
00 06 00 04 00 03 00 19 00 03 00 0c 00 00 00 00 00 00 00
00 00 00 00 00 19 00 0c 00 00 00 00 00 00 00 00 00 00 00
```

show subscribers (LNS on MX Series Routers)

user@host> show subscribers

Interface	IP Address/VLAN ID	User Name	LS:RI
si-4/0/0.1	192.0.2.0	user@example.com	default:default

show subscribers (L2TP Switched Tunnels)

user@host> show subscribers

Interface	IP Address/VLAN ID	User Name	LS:RI
si-2/1/0.1073741842	Tunnel-switched	user@example.com	default:default
si-2/1/0.1073741843	Tunnel-switched	user@example.com	default:default

show subscribers aggregation-interface-set-name

user@host> show subscribers aggregation-interface-set-name FRA*

Interface	IP Address/VLAN ID	User Name
LS:RI		
ge-1/0/0.3221225472	50	anyp
default:ispl-subscriber		

show subscribers client-type dhcp detail

user@host> show subscribers client-type dhcp detail

```
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
```

```

Interface type: Dynamic
Dynamic Profile Name: dhcp-demux
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: user :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 203.0.113.27
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:f3
State: Active
Radius Accounting ID: 1234 :2560
Login Time: 2009-08-25 14:43:56 PDT

```

show subscribers client-type dhcp detail (DHCPv6)

user@host> **show subscribers client-type dhcp detail**

```

Type: DHCP
User Name: DEFAULTUSER
IPv6 Address: 2001:db8::2
IPv6 Prefix: 2001:db8:1::/64
Logical System: default
Routing Instance: default
Interface: demux0.3221225602
Interface type: Static
Underlying Interface: demux0.3221225602
Dynamic Profile Name: client-profile
MAC Address: 00:00:5E:00:53:01
State: Active
Radius Accounting ID: 142
Session ID: 142
PFE Flow ID: 148
Stacked VLAN Id: 1
VLAN Id: 1
Login Time: 2018-03-29 12:27:38 EDT
DHCP Options: len 56
00 08 00 02 00 00 00 01 00 0e 00 01 00 01 22 4f d0 33 00 11

```

```

01 00 00 01 00 03 00 0c 00 00 00 0a 00 04 9d 40 00 07 62 00
00 19 00 0c 00 00 00 0b 00 04 9d 40 00 07 62 00
Server DHCPV6 Options: len 94
00 0a 00 06 11 22 33 44 55 66 00 11 00 09 00 00 0c 4c 00 02
00 01 aa 00 11 00 20 00 00 0a 4c 00 02 00 02 32 33 00 03 00
03 34 35 36 00 05 00 06 31 32 33 34 35 36 00 06 00 01 31 00
11 00 09 00 00 0b 4c 00 02 00 01 bb 00 11 00 12 00 00 0d e9
00 01 00 03 aa bb cc 00 02 00 03 dd ee cc
DHCPV6 Header: len 4
01 fc e4 96

```

show subscribers client-type dhcp extensive

user@host> show subscribers client-type dhcp extensive

```

Type: DHCP
User Name: user
IP Address: 192.0.2.4
IP Netmask: 255.0.0.0
IPv6 Address: 2001:db8:3::103
IPv6 Prefix: 2001:db8::/68
Domain name server inet6: 2001:db8:1 abcd::2
Logical System: default
Routing Instance: default
Interface: ge-0/0/0.0
Interface type: Static
Underlying Interface: ge-0/0/0.0
MAC Address: 00:00:5e:00:53:01
State: Configured
Radius Accounting ID: 10
Session ID: 10
PFE Flow ID: 2
VLAN Id: 100
Agent Circuit ID: ge-0/0/0:100
Agent Remote ID: ge-0/0/0:100
Login Time: 2017-05-23 12:52:22 IST
DHCPV6 Options: len 69
00 01 00 0e 00 01 00 01 59 23 e3 31 00 10 94 00 00 01 00 08
00 02 00 00 00 19 00 29 00 00 00 00 00 04 9d 40 00 07 62 00
00 1a 00 19 00 09 3a 80 00 27 8d 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
Server DHCP Options: len 13
3a 04 00 00 00 ff 00 3b 04 00 00 0f 00
Server DHCPV6 Options: len 8

```



```

00 0a 00 04 ab cd ef ab
DHCPV6 Header: len 4
01 00 00 04
IP Address Pool: al_pool30
IPv6 Address Pool: ia_na_pool
IPv6 Delegated Address Pool: prefix_delegate_pool

```

show subscribers client-type fixed-wireless-access

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

Interface	IP Address/VLAN ID	User Name
LS:RI		
ps1.3221225472	192.0.2.10	505024101215074
default:default		
ps1.3221225473	192.0.2.11	505024101215075
default:default		

show subscribers client-type fixed-wireless-access detail (Detail)

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

```

Type: FWA
User Name: 505024101215074
IP Address: 192.0.2.10
IP Netmask: 255.255.0.0
Interface: ps1.3221225472
Interface type: Dynamic
Dynamic Profile Name: fwa-profile
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 11
Login Time: 2019-04-10 14:10:12 PDT
Local TEID-U: 1
Local TEID-C: 1
Remote TEID-U: 2000000
Remote TEID-C: 1000000
GTPU Tunnel Remote IP Address: 203.0.113.1.3
GTPU Tunnel Local IP Address: 203.0.113.2.5
GTPC Remote IP Address: 203.0.113.1.2

```

GTPC Local IP Address: 203.0.113.1.1
Access Point Name: user21

show subscribers client-type hag-bundle

user@host> show subscribers client-type hag-bundle

Interface	IP Address/VLAN ID	User Name
LS:RI		
ps0.3221225473		user1
default:default		

show subscribers client-type hag-bundle (Detail)

user@host> show subscribers client-type hag-bundle detail

```
Type: HAG-BUNDLE
User Name: user1
Logical System: default
Routing Instance: default
Interface: ps0.3221225473
Interface type: Dynamic
Dynamic Profile Name: hag-dyn-profile
State: Active
Radius Accounting ID: 2
Session ID: 2
HAG Session ID: 1634170802
PFE Flow ID: 9
Login Time: 2019-08-13 16:36:02 EDT
```

show subscribers client-type -hag-tunnel

user@host> show subscribers client-type hag-tunnel

Interface	IP Address/VLAN ID	User Name
LS:RI		
ps0.3221225472		user1
default:default		
ps0.3221225473		user1
default:default		

show subscribers client-type hag-tunnel (Detail)

user@host> **show subscribers client-type hag-tunnel detail**

```
Type: HAG-TUNNEL
User Name: user1
Logical System: default
Routing Instance: default
Interface: ps0.3221225472
Interface type: Dynamic
Dynamic Profile Name: hag-access-line-profile
State: Active
Session ID: 1
Bundle Session ID: 2
HAG Session ID: 1634170802
HAG Tunnel Type: LTE
PFE Flow ID: 7
Login Time: 2019-08-13 16:36:02 EDT
```

```
Type: HAG-TUNNEL
User Name: user1
Logical System: default
Routing Instance: default
Interface: ps0.3221225474
Interface type: Dynamic
Dynamic Profile Name: hag-access-line-profile
State: Active
Session ID: 3
Bundle Session ID: 2
HAG Session ID: 1634170802
HAG Tunnel Type: DSL
PFE Flow ID: 10
Login Time: 2019-08-13 16:36:02 EDT
```

show subscribers client-type vlan-oob detail

user@host> **show subscribers client-type vlan-oob detail**

```
Type: VLAN-OOB
User Name: L2WS.line-aci-1.line-ari-1
Logical System: default
Routing Instance: ISPl
Interface: demux0.1073744127
Interface type: Dynamic
Underlying Interface: ge-1/0/0
```

```

Dynamic Profile Name: Prof_L2WS
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 1234
Session ID: 77
VLAN Id: 126
Core-Facing Interface: ge-2/1/1
VLAN Map Id: 6
Inner VLAN Map Id: 2001
Agent Circuit ID: line-aci-1
Agent Remote ID: line-ari-1
Login Time: 2013-10-29 14:43:52 EDT

```

show subscribers count

```
user@host> show subscribers count
```

```
Total Subscribers: 188, Active Subscribers: 188
```

show subscribers address detail (IPv6)

```
user@host> show subscribers address 203.0.113.137 detail
```

```

Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 203.0.113.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 2001:db8:0:c88::/32
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:00:5e:00:53:53
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: example demux0.8201:6544
Session ID: 6544
Agent Circuit ID: ifl3720
Agent Remote ID: ifl3720

```

```
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1
```

show subscribers detail (IPv4)

```
user@host> show subscribers detail
```

```
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Primary DNS Address: 192.0.2.0
Secondary DNS Address: 192.0.2.1
Primary WINS Address: 192.0.2.3
Secondary WINS Address: 192.0.2.4
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: example :2304
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2
```

show subscribers detail (IPv6)

```
user@host> show subscribers detail
```

```
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8:ffff:1::/32
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03
State: Active
```

```

Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00

```

show subscribers detail (pseudowire Interface for GRE Tunnel)

user@host> show subscribers detail

Interface	IP Address/VLAN ID	User Name	LS:RI
ps0.3221225484	192.0.2.2		
ps0.3221225485	192.0.2.3		
demux0.3221225486	1		default:default
demux0.3221225487	1		default:default
demux0.3221225488	198.51.0.1		default:default
demux0.3221225489	198.51.0.2		default:default

show subscribers detail (IPv6 Static Demux Interface)

user@host> show subscribers detail

```

Type: STATIC-INTERFACE
User Name: user@example.com
IPv6 Prefix: 2001:db8:3:4:5:6:7:aa/32
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT

```

show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

user@host> show subscribers detail

```

Type: L2TP
User Name: user@example.com
IP Address: 203.0.113.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST

```

show subscribers detail (L2TP Switched Tunnels)

user@host> **show subscribers detail**

```

Type: L2TP
User Name: user@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 203.0.113.51
Remote IP Address: 192.0.2.0
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST

```

```

Type: L2TP
User Name: user@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741843
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched

```

```
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 203.0.113.31
Remote IP Address: 192.0.2.1
Session ID: 22
Login Time: 2013-01-18 03:01:14 PST
```

show subscribers detail (Tunneled Subscriber)

```
user@host> show subscribers detail
```

```
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512
```

show subscribers detail (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers detail
```

```
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST
```

```
Type: PPPoE
User Name: dualstackuser1@example1.com
IP Address: 203.0.113.13
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
```



```

Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00

```

show subscribers detail (ACI Interface Set Session)

user@host> show subscribers detail

```

Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT

```

show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```
user@host> show subscribers detail
```

```
Type: PPPoE
User Name: ppphint2
IP Address: 203.0.113.15
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST
```

show subscribers extensive

```
user@host> show subscribers extensive
```

```
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8:ffff:1::/32
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
```

```
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48
```

show subscribers extensive (Aggregation Node Interface Set and DSL Forum Attributes)

```
user@host> show subscribers extensive
```

```
Type: VLAN-OOB
User Name: ancp
Logical System: default
Routing Instance: ispl-subscriber
Interface: ge-1/0/0.3221225472
Interface type: Dynamic
Interface Set: FRA-DPU-C-100
Underlying Interface: ge-1/0/0
Core IFL Name: ge-1/0/4.0
Dynamic Profile Name: Prof_L2BSA
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 13
VLAN Id: 50
VLAN Map Id: 20
Inner VLAN Map Id: 1
Inner VLAN Tag Protocol Id: 0x88a8
Agent Circuit ID: circuit 201
Agent Remote ID: remote-id
Aggregation Interface-set Name: FRA-DPU-C-100
Login Time: 2018-05-29 08:43:42 EDT
Accounting interval: 72000
Dynamic configuration:
  junos-cos-scheduler-map: 100m
  junos-inner-vlan-tag-protocol-id: 0x88a8
  junos-vlan-map-id: 20

Type: PPPoE
IP Address: 192.85.128.1
IP Netmask: 255.255.255.255
Logical System: default
Routing Instance: default
Interface: pp0.3221225474
Interface type: Dynamic
Interface Set: ge-1/0/0
Underlying Interface: demux0.3221225473
```

```

Dynamic Profile Name: pppoe-client-profile-with-cos
MAC Address: 00:10:94:00:00:03
State: Active
Radius Accounting ID: 3
Session ID: 3
PFE Flow ID: 16
Stacked VLAN Id: 50
VLAN Id: 7
Agent Circuit ID: circuit 201
Agent Remote ID: remote-id
Aggregation Interface-set Name: FRA-DPU-C-100
Login Time: 2018-05-29 08:43:45 EDT
IP Address Pool: pool-1
Accounting interval: 72000
DSL type: G.fast
Frame/cell mode: Frame
Overhead accounting bytes: 10
Actual upstream data rate: 100000 kbps
Actual downstream data rate: 200000 kbps
Calculated downstream data rate: 180000 kbps
Calculated upstream data rate: 90000 kbps
Adjusted upstream data rate: 80000 kbps
Adjusted downstream data rate: 160000 kbps
DSL Line Attributes
  Agent Circuit ID: circuit 201
  Agent Remote ID: remote-id
  Actual upstream data rate: 100000
  Actual downstream data rate: 200000
  DSL type: G.fast
  Access Aggregation Circuit ID: #FRA-DPU-C-100
  Attribute type: 0xAA, Attribute length: 4
    198 51 100 78

```

show subscribers extensive (Passive Optical Network Circuit Interface Set)

user@host> show subscribers client-type dhcp extensive

```

Type: DHCP
IP Address: 192.0.2.136
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073741842
Interface type: Dynamic

```

```

Interface Set: ot101.xyz101-202
Underlying Interface: demux0.1073741841
Dynamic Profile Name: dhcp-profile
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: user :19
Session ID: 19
VLAN Id: 1100
Agent Remote ID: ABCD01234|100M|AAAA01234|ot101.xyz101-202

Login Time: 2017-03-29 10:30:46 PDT
DHCP Options: len 97
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 02 33 04 00 00
17 70 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
32 2d 31 2d 31 37 05 01 06 0f 21 2c 52 2b 02 29 41 42 43 44
30 31 32 33 34 7c 31 30 30 4d 7c 41 41 41 41 30 31 32 33 34
7c 6f 74 6c 30 31 2e 78 79 7a 31 30 31 2d 32 30 32
IP Address Pool: POOL-V4

```

show subscribers extensive (DNS Addresses from Access Profile or Global Configuration)

user@host> show subscribers extensive

```

Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Domain name server inet: 198.51.100.1 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f

```

```
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

show subscribers extensive (DNS Addresses from RADIUS)

```
user@host> show subscribers extensive
```

```
Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
IPv6 Primary DNS Address: 2001:db8:5001::12
IPv6 Secondary DNS Address: 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool
```

show subscribers extensive (IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration)

```
user@host> show subscribers extensive
```

```
Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
```

```

IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool

```

show subscribers extensive (RPF Check Fail Filter)

user@host> show subscribers extensive

```

...
Type: VLAN
  Logical System: default
  Routing Instance: default
  Interface: ae0.1073741824
  Interface type: Dynamic
  Dynamic Profile Name: vlan-prof
  State: Active
  Session ID: 9
  VLAN Id: 100
  Login Time: 2011-08-26 08:17:00 PDT
  IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
  IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...

```

show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

user@host> show subscribers extensive

```

Type: L2TP
User Name: user@example.com

```

```

IP Address: 203.0.113.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out

```

show subscribers extensive (IPv4 and IPv6 Dual Stack)

user@host> show subscribers extensive

```

Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@example1.com
IP Address: 203.0.113.13
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2

```



```

Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2001:db8:2016:1:1::1/64
IPv6 Framed Interface Id: 1:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

```

```

Type: DHCP
IPv6 Prefix: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48

```

show subscribers extensive (ADF Rules)

user@host> show subscribers extensive

```

...
Service Session ID: 12
Service Session Name: SERVICE-PROFILE
State: Active
Family: inet
  ADF IPv4 Input Filter Name: __junos_adf_12-demux0.3221225474-inet-in
    Rule 0: 010101000b0101020b020200201811
      from {
        source-address 203.0.113.232;
        destination-address 198.51.100.0/24;
        protocol 17;
      }

```

```

    }
    then {
        accept;
    }

```

show subscribers extensive (Effective Shaping-Rate)

user@host> show subscribers extensive

```

Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741837
Interface type: Dynamic
Interface Set: ifset-1
Underlying Interface: ael
Dynamic Profile Name: svlan-dhcp-test
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.201
VLAN Id: 0x8100.201
Login Time: 2011-11-30 00:18:04 PST
Effective shaping-rate: 31000000k
...

```

show subscribers extensive (PPPoE Subscriber Access Line Rates)

user@host> show subscribers extensive

```

Type: PPPoE
IP Address: 198.51.100.1
IP Netmask: 255.255.255.255
Logical System: default
Routing Instance: default
Interface: pp0.3221225475
Interface type: Dynamic
Underlying Interface: demux0.3221225474
Dynamic Profile Name: pppoe-client-profile-with-cos
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 4
Session ID: 4
PFE Flow ID: 14

```

Stacked VLAN Id: 40
 VLAN Id: 1
 Agent Circuit ID: circuit0
 Agent Remote ID: remote0
 Login Time: 2017-04-06 15:52:32 PDT

User Name: DAVE-L2BSA-SERVICE
 Logical System: default
 Routing Instance: isp-1-subscriber
 Interface: ge-1/2/4.3221225472
 Interface type: Dynamic
 Interface Set: ge-1/2/4
 Underlying Interface: ge-1/2/4
 Core IFL Name: ge-1/3/4.0
 Dynamic Profile Name: L2BSA-88a8-400LL1300VO
 State: Active

Radius Accounting ID: 1

Session ID: 1
 PFE Flow ID: 14
 VLAN Id: 13
 VLAN Map Id: 102
 Inner VLAN Map Id: 1
 Agent Circuit ID: circuit-aci-3
 Agent Remote ID: remote49-3
 Login Time: 2017-04-05 16:59:29 EDT
 Service Sessions: 4
 IFL Input Filter Name: L2BSA-CP-400LL1300VO-ge-1/2/4.3221225472-in
 IFL Output Filter Name: L2BSA-CP-400LL1300VO-ge-1/2/4.3221225472-out
 Accounting interval: 900

DSL type: VDSL

Frame/Cell Mode: Frame

Overhead accounting bytes: -10

Actual upstream data rate: 1024 kbps

Actual downstream data rate: 4096 kbps

Adjusted downstream data rate: 3686 kbps

Adjusted upstream data rate: 922 kbps

Dynamic configuration:

junos-vlan-map-id: 102

Service Session ID: 5

Service Session Name: SRL-L1

State: Active

Family: inet, inet6

IFL Input Filter Name: L2BSA-FWF-in-10048-ge-1/2/4.3221225472-in

IFL Output Filter Name: L2BSA-FWF-out-25088-ge-1/2/4.3221225472-out

```

Service Activation time: 2017-04-05 16:59:30 EDT
Dynamic configuration:
  l2bsa-fwf-in: L2BSA-FWF-in-10048
  l2bsa-fwf-out: L2BSA-FWF-out-25088
  rldown: 25088
  rlup: 10048

```

show subscribers extensive (Subscriber Session Using PCEF Profile)

user@host> **show subscribers extensive**

```

Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221225517
Interface type: Dynamic
Underlying Interface: ge-1/0/3
Dynamic Profile Name: svlan-dhcp
State: Active
Session ID: 59
PFE Flow ID: 71
Stacked VLAN Id: 0x8100.1
VLAN Id: 0x8100.2
Login Time: 2017-03-28 08:23:08 PDT

Type: DHCP
User Name: pcefuser
IP Address: 192.0.2.26
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225518
Interface type: Dynamic
Underlying Interface: demux0.3221225517
Dynamic Profile Name: dhcp-client-prof
MAC Address: 00:00:5e:00:53:01
State: Active
Radius Accounting ID: 60
Session ID: 60
PFE Flow ID: 73
Stacked VLAN Id: 1
VLAN Id: 2
Login Time: 2017-03-28 08:23:08 PDT
Service Sessions: 1

```

```

DHCP Options: len 9
35 01 01 37 04 01 03 3a 3b
IP Address Pool: pool-ipv4
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
PCEF Profile: pcef-prof-1
PCEF Rule/Rulebase: default
Dynamic configuration:
  junos-input-service-filter: svc-filt-1
  junos-input-service-set: tdf-service-set
  junos-output-service-filter: svc-filt-1
  junos-output-service-set: tdf-service-set
  junos-pcef-profile: pcef-prof-1
  junos-pcef-rule: default

Service Session ID: 61
Service Session Name: pcef-serv-prof
State: Active
Family: inet
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
PCEF Profile: pcef-prof-1
PCEF Rule/Rulebase: limit-fb
Service Activation time: 2017-03-28 08:31:19 PDT
Dynamic configuration:
  pcef-prof: pcef-prof-1
  pcef-rule1: limit-fb
  svc-filt: svc-filt-1
  svc-set: tdf-service-set

```

show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

user@host> **show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail**

```

Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10

```

```

Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 203.0.113.17
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address:
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT

```

show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```
user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
```

```

Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 203.0.113.17
Logical System: default
Routing Instance: default
Interface: pp0.1073741834

```

```

Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:5e:00:53:52
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT

```

show subscribers id accounting-statistics

```
user@host> show subscribers id 601 accounting-statistics
```

```

Session ID: 601
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

show subscribers interface accounting-statistics

```
user@host> show subscribers interface pp0.3221226949 accounting-statistics
```

```

Session ID: 501
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
Output bytes : 0

```

```

Input packets: 0
Output packets: 0

Session ID: 502
Accounting Statistics:
Input bytes : 87654
Output bytes : 72108
Input packets: 3322
Output packets: 3322
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

```

Session ID: 503
Accounting Statistics:
Input bytes : 156528
Output bytes : 123865
Input packets: 7448
Output packets: 7448
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

show subscribers interface extensive

user@host> show subscribers interface demux0.1073741826 extensive

```

Type: VLAN
User Name: user@test.example.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:5e:00:53:04
State: Active
Radius Accounting ID: 12
Session ID: 12
Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902

```



```
Login Time: 2011-10-20 16:21:59 EST
```

```
Type: DHCP
```

```
User Name: user@test.example.com
```

```
IP Address: 192.0.2.0
```

```
IP Netmask: 255.255.255.0
```

```
Logical System: default
```

```
Routing Instance: testnet
```

```
Interface: demux0.1073741826
```

```
Interface type: Static
```

```
MAC Address: 00:00:5e:00:53:04
```

```
State: Active
```

```
Radius Accounting ID: 21
```

```
Session ID: 21
```

```
Login Time: 2011-10-20 16:24:33 EST
```

```
Service Sessions: 2
```

```
Service Session ID: 25
```

```
Service Session Name: SUB-QOS
```

```
State: Active
```

```
Service Session ID: 26
```

```
Service Session Name: service-cb-content
```

```
State: Active
```

```
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
```

```
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out
```

show subscribers logical-system terse

```
user@host> show subscribers logical-system test1 terse
```

Interface	IP Address/VLAN ID	User Name	LS:RI
demux0.1073741825	203.0.113.3	RETAILER1-CLIENT	test1:retailer1
demux0.1073741826	203.0.113.4	RETAILER2-CLIENT	test1:retailer2

show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
```

```
Total subscribers: 3998, Active Subscribers: 3998
```

show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
```

```
Total Subscribers: 188, Active Subscribers: 183
```

show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
```

```
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
```

```
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
```

```
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
```

```
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers user-name detail

user@host> show subscribers user-name larry1 detail

```
Type: DHCP
User Name: larry1
IP Address: 203.0.113.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
MAC Address: 00:00:5e:00:53:01
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c
```

show subscribers vlan-id

user@host> show subscribers vlan-id 100

Interface	IP Address	User Name
ge-1/0/0.1073741824		
ge-1/2/0.1073741825		

show subscribers vlan-id detail

user@host> show subscribers vlan-id 100 detail

```
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
```

```

State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

```

show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session)

user@host> show subscribers vpi 40 vci 50 extensive

```

Type: PPPoE
User Name: testuser
IP Address: 203.0.113.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
Login Time: 2012-12-03 07:49:26 PST
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102

```

show subscribers address detail (Enhanced Subscriber Management)

user@host> show subscribers address 203.0.113.111 detail

```

Type: DHCP
User Name: simple_filters_service
IP Address: 203.0.113.111
IP Netmask: 255.0.0.0
Logical System: default

```

```
Routing Instance: default
Interface: demux0.3221225482
Interface type: Dynamic
Underlying Interface: demux0.3221225472
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:0f
State: Active
Radius Accounting ID: 11
Session ID: 11
PFE Flow ID: 15
Stacked VLAN Id: 210
VLAN Id: 209
Login Time: 2014-03-24 12:53:48 PDT
Service Sessions: 1
DHCP Options: len 3
35 01 01
```

show subscribers summary

Syntax

```
show subscribers summary
<all>
<detail | extensive | terse>
<count>
<physical-interface physical-interface-name>
<logical-system logical-system pic | port | routing-instance routing-instance | slot>
```

Release Information

Command introduced in Junos OS Release 10.2.

Description

Display summary information for subscribers.

Options

none—Display summary information by state and client type for all subscribers.

all—(Optional) Display summary information by state, client type, and LS:RI.

detail | extensive | terse—(Not supported on MX Series routers) (Optional) Display the specified level of output.

count—(Not supported on MX Series routers) (Optional) Display the count of total subscribers and active subscribers for any specified option.

logical-system *logical-system*—(Optional) Display subscribers whose logical system matches the specified logical system.

physical-interface *physical-interface-name*—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers whose physical interface matches the specified physical interface, by subscriber state, client type, and LS:RI.

pic—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by PIC number and the total number of subscribers.

port—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by port number and the total number of subscribers.

routing-instance *routing-instance*—(Optional) Display subscribers whose routing instance matches the specified routing instance.

slot—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by FPC slot number and the total number of subscribers.

NOTE: Due to display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

RELATED DOCUMENTATION

[show subscribers](#) | [1527](#)

List of Sample Output

[show subscribers summary on page 1579](#)

[show subscribers summary all on page 1580](#)

[show subscribers summary physical-interface on page 1580](#)

[show subscribers summary physical-interface pic on page 1581](#)

[show subscribers summary physical-interface port on page 1581](#)

[show subscribers summary physical-interface slot on page 1582](#)

[show subscribers summary pic on page 1582](#)

[show subscribers summary pic \(Aggregated Ethernet Interfaces\) on page 1582](#)

[show subscribers summary port on page 1582](#)

[show subscribers summary port \(Pseudowire Interfaces\) on page 1583](#)

[show subscribers summary port extensive on page 1583](#)

[show subscribers summary slot on page 1583](#)

[show subscribers summary terse on page 1584](#)

Output Fields

[Table 95 on page 1578](#) lists the output fields for the **show subscribers summary** command. Output fields are listed in the approximate order in which they appear.

Table 95: show subscribers summary Output Fields

Field Name	Field Description	Level of Output
Subscribers by State	<p>Number of subscribers summarized by state. The summary information includes the following:</p> <ul style="list-style-type: none"> • Init—Number of subscriber currently in the initialization state. • Configured—Number of configured subscribers. • Active—Number of active subscribers. • Terminating—Number of subscribers currently terminating. • Terminated—Number of terminated subscribers. • Total—Total number of subscribers for all states. 	detail none
Subscribers by Client Type	<p>Number of subscribers summarized by client type. Client types can include DHCP, GRE, HAG-BUNDLE, HAG-TUNNEL, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN, and VLAN-OOB. Also displays the total number of subscribers for all client types (Total).</p>	detail extensive none
Subscribers by LS:RI	<p>Number of subscribers summarized by logical system:routing instance (LS:RI) combination. Also displays the total number of subscribers for all LS:RI combinations (Total).</p>	detail none
Subscribers by Connection Type	<p>Number of subscribers summarized by connection type, Cross-connected or Terminated.</p>	extensive
Interface	<p>Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface.</p> <p>The * character indicates a continuation of addresses for the same session.</p> <p>For aggregated Ethernet interfaces, the output of the summary (pic port slot) options prefixes the interface name with ae0:.</p> <p>For pseudowire IFDs, this field displays both the pseudowire and the associated logical tunnel (LT) and redundant logical tunnel (RLT) anchor interface. For example:</p> <pre>ps0: lt-2/1/0 ps1: rlt0: lt-4/0/0</pre>	All levels

Table 95: show subscribers summary Output Fields (*continued*)

Field Name	Field Description	Level of Output
Count	<p>Count of subscribers displayed for each PIC, port, or slot when those options are specified with the summary option. For an aggregated Ethernet configuration, the total subscriber count does not equal the sum of the individual PIC, port, or slot counts, because each subscriber can be in more than one aggregated Ethernet link.</p> <p>Multiple pseudowire interfaces can share a given logical tunnel or redundant logical tunnel anchor interface. Starting in Junos OS Release 18.1R1, the field displays subscribers per individual pseudowire interface.</p> <p>In earlier releases, the field displays the same number of subscribers for all pseudowire interfaces that share the same tunnel interface as their anchor point.</p>	detail extensive none
Total Subscribers	Total number of subscribers for all physical interfaces, all PICs, all ports, or all LS:RI slots.	detail extensive none
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i>	terse
User Name	Name of subscriber.	terse
LS:RI	Logical system and routing instance associated with the subscriber.	terse

Sample Output

show subscribers summary

user@host> **show subscribers summary**

```
Subscribers by State
Init          3
Configured    2
Active       188
Terminating   2
Terminated    1

TOTAL        191
```

Subscribers by Client Type

DHCP	107
HAG-TUNNEL	2
HAG-BUNDLE	1
PPP	76
VLAN	8
VLAN-OOB	2
TOTAL	196

show subscribers summary all

```
user@host> show subscribers summary all
```

Subscribers by State

Init	3
Configured	2
Active	183
Terminating	2
Terminated	1
 TOTAL	 191

Subscribers by Client Type

DHCP	107
PPP	76
VLAN	8
 TOTAL	 191

Subscribers by LS:RI

default:default	1
default:ri1	28
default:ri2	16
ls1:default	22
ls1:riA	38
ls1:riB	44
logsysX:routinstY	42
 TOTAL	 191

show subscribers summary physical-interface

```
user@host> show subscribers summary physical-interface ge-1/0/0
```

```
Subscribers by State
  Active: 3998
  Total: 3998

Subscribers by Client Type
  DHCP: 3998
  Total: 3998

Subscribers by LS:RI
  default:default: 3998
  Total: 3998
```

show subscribers summary physical-interface pic

user@host> **show subscribers summary physical-interface ge-0/2/0 pic**

```
Subscribers by State
  Active: 4825
  Total: 4825

Subscribers by Client Type
  DHCP: 4825
  Total: 4825

Subscribers by LS:RI
  default:default: 4825
  Total: 4825
```

show subscribers summary physical-interface port

user@host> **show subscribers summary physical-interface ge-0/3/0 port**

```
Subscribers by State
  Active: 4825
  Total: 4825

Subscribers by Client Type
  DHCP: 4825
  Total: 4825

Subscribers by LS:RI
  default:default: 4825
  Total: 4825
```

show subscribers summary physical-interface slot

```
user@host> show subscribers summary physical-interface ge-2/0/0 slot
```

```
Subscribers by State
  Active: 4825
  Total: 4825

Subscribers by Client Type
  DHCP: 4825
  Total: 4825

Subscribers by LS:RI
  default:default: 4825
  Total: 4825
```

show subscribers summary pic

```
user@host> show subscribers summary pic
```

```
Interface      Count
ge-1/0         1000
ge-1/3         1000

Total Subscribers: 2000
```

show subscribers summary pic (Aggregated Ethernet Interfaces)

```
user@host> show subscribers summary pic
```

```
Interface      Count
ae0: ge-1/0    801
ae0: ge-1/3    801

Total Subscribers: 801
```

show subscribers summary port

```
user@host> show subscribers summary port
```

```
Interface      Count
ge-5/0/1       201
ge-5/0/2       301
```

```
Total Subscribers: 502
```

show subscribers summary port (Pseudowire Interfaces)

```
user@host> show subscribers summary port
```

```
ps0: lt-2/1/0 10
ps1: lt-2/1/0 20

Total Subscribers: 30
```

show subscribers summary port extensive

```
user@host>show subscribers summary port extensive
```

```
Interface: ge-5/0/1
Count: 201
Detail:
Subscribers by Client Type
  DHCP: 100
  PPPoE: 100
  VLAN-OOB: 1
Subscribers by Connection Type
  Terminated: 200
  Cross-connected: 1

Interface: ge-5/0/2
Count: 301
Detail:
Subscribers by Client Type
  DHCP: 200
  PPPoE: 100
  VLAN-OOB: 1
Subscribers by Connection Type
  Terminated: 300
  Cross-connected: 1

Total Subscribers: 502
```

show subscribers summary slot

```
user@host> show subscribers summary slot
```

Interface	Count
ge-1	2000

Total Subscribers: 2000

show subscribers summary terse

user@host> show subscribers summary terse

Interface	IP Address/VLAN ID	User Name	LS:RI
ge-1/3/0.1073741824	100		default:default
demux0.1073741824	203.0.113.10	WHOLESALE-CLIENT	default:default
demux0.1073741825	203.0.113.13	RETAILER1-CLIENT	test1:retailer1
demux0.1073741826	203.0.113.213	RETAILER2-CLIENT	test1:retailer2

show tcp-forwarding status

Syntax

```
show tcp-forwarding status
listening-port port-number listening-address ipv4-listening-address
routing-instance routing-instance-name
```

Release Information

Command introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Display the status of TCP mapping and the current TCP connections for each mapping. You can limit the display to a specific listening port/listening address combination, per routing instance. If you do not specify a routing instance, the default routing instance is assumed.

Options

listening-address *ipv4-listening-address*—IPv4 address that is part of a listening port/listening address combination. The listening address is one on the BNG that external management systems or remote devices must use when attempting to trigger connections on the listening port. You must also specify a listening port.

listening-port *port-number*—Port number that is part of a listening port/listening address combination. The listening port is one that the BNG monitors for connections to be triggered by external management systems or remote devices.

Range: 8000 through 8031

routing-instance *routing-instance-name*—Name of the routing instance for the TCP mapping.

Required Privilege Level

view

RELATED DOCUMENTATION

[TCP Port Forwarding for Remote Device Management | 655](#)

List of Sample Output

[show tcp-forwarding status on page 1587](#)

[show tcp-forwarding status \(Listening Port and Address\) on page 1589](#)

Output Fields

[Table 96 on page 1586](#) lists the output fields for the **show tcp-forwarding status** command. Output fields are listed in the approximate order in which they appear.

Table 96: show tcp-forwarding status Output Fields

Field Name	Field Description
Listening on	<p>Routing instance, IPv4 address, and port that the BNG is monitoring for connection triggers from external management systems or remote devices, in the following format:</p> <p>[routing-instance-name]:ip-address:port-number</p> <p>The following status and statistics are displayed for the TCP listening connection:</p> <ul style="list-style-type: none"> • Status—Current state of the BNG regarding the TCP connection: <ul style="list-style-type: none"> • listening—The connection is active. • not-listening—The connection is inactive. This is typically indicative of a misconfiguration, such as an invalid listening address for the routing instance. • Total Bytes—Total number of bytes sent and received on the TCP listening connection. <ul style="list-style-type: none"> • Rx—Received bytes. • Tx—Transmitted bytes.
Forwarding to	<p>Routing instance, IPv4 address, and port number where the BNG is forwarding data on the TCP connection, in the following format:</p> <p>[routing-instance-name]:ip-address:port-number</p> <p>The following statistics are displayed for the TCP port forwarding connection:</p> <ul style="list-style-type: none"> • Total Bytes—Total number of bytes sent and received on the TCP port forwarding connection. <ul style="list-style-type: none"> • Rx—Received bytes. • Tx—Transmitted bytes.
Allowed Source Prefixes	<p>IPv4 source prefixes from which the BNG can accept connection requests. Request from all other addresses are rejected. A /32 mask indicates a single acceptable address.</p> <p>The allowed values are compared to the source address in the TCP header from the triggering entity.</p>

Table 96: show tcp-forwarding status Output Fields (*continued*)

Field Name	Field Description
Connections	<p>Status and statistics for the configured TCP connections.</p> <ul style="list-style-type: none"> • Max—Configured maximum number of connections allowed on the listening port. • Active—Number of TCP connections that are currently active on the listening port. • Source—List of source address and source port combinations with which the BNG has a current connection. The field value is displayed in the following format: <i>source-ip-address:port-number</i> • Listening and Forwarding—Status of the listening and forwarding connections. <ul style="list-style-type: none"> • connected—The connection is active. • disconnected—The connection is removed. This is a transient state, because the TCP connection pair is removed when disconnected. • Bytes—Number of bytes transmitted and received on the listening or forwarding connections for the source address/port combination. <ul style="list-style-type: none"> • Rx—Received bytes. • Tx—Transmitted bytes.

Sample Output

show tcp-forwarding status

user@host> **show tcp-forwarding status**

```

Listening on: [default:]203.0.113.50:8000
  Status: listening
  Total Bytes Rx: 0 Tx: 0
Forwarding to: [default:]192.0.0.2:830
  Total Bytes Rx: 0 Tx: 0
Allowed Source Prefixes:
  198.51.100.3/32
Connections Max: 1 Active: 1
  Source: 198.51.100.3:55000
  Listening: connected Bytes Rx: 0 Tx: 0

```

```
Forwarding: connected Bytes Rx: 0 Tx: 0

Listening on: [default:]203.0.113.50:8001
  Status: listening
  Total Bytes Rx: 0 Tx: 0
Forwarding to: [default:]192.0.0.3:830
  Total Bytes Rx: 0 Tx: 0
Allowed Source Prefixes:
  198.51.100.3/32
Connections Max: 1 Active: 1
Source: 198.51.100.3:55001
  Listening: connected Bytes Rx: 0 Tx: 0
  Forwarding: connected Bytes Rx: 0 Tx: 0

Listening on: [default:]203.0.113.50:8002
  Status: listening
  Total Bytes Rx: 0 Tx: 0
Forwarding to: [default:]192.0.0.4:830
  Total Bytes Rx: 0 Tx: 0
Allowed Source Prefixes:
  198.51.100.3/32
Connections Max: 1 Active: 1
Source: 198.51.100.3:55002
  Listening: connected Bytes Rx: 0 Tx: 0
  Forwarding: connected Bytes Rx: 0 Tx: 0

Listening on: [default:]203.0.113.50:8003
  Status: listening
  Total Bytes Rx: 0 Tx: 0
Forwarding to: [default:]192.0.0.5:830
  Total Bytes Rx: 0 Tx: 0
Allowed Source Prefixes:
  198.51.100.3/32
Connections Max: 1 Active: 1
Source: 198.51.100.3:55003
  Listening: connected Bytes Rx: 0 Tx: 0
  Forwarding: connected Bytes Rx: 0 Tx: 0

Listening on: [default:]203.0.113.50:8020
  Status: listening
  Total Bytes Rx: 0 Tx: 0
Forwarding to: [default:]198.51.100.1:49
  Total Bytes Rx: 0 Tx: 0
Allowed Source Prefixes:
```

```

192.0.0.1/24
Connections Max: 4 Active: 4
Source: 192.0.0.2:55000
  Listening: connected Bytes Rx: 0 Tx: 0
  Forwarding: connected Bytes Rx: 0 Tx: 0
Source: 192.0.0.3:55000
  Listening: connected Bytes Rx: 0 Tx: 0
  Forwarding: connected Bytes Rx: 0 Tx: 0
Source: 192.0.0.4:56022
  Listening: connected Bytes Rx: 0 Tx: 0
  Forwarding: connected Bytes Rx: 0 Tx: 0
Source: 192.0.0.5:55000
  Listening: connected Bytes Rx: 0 Tx: 0
  Forwarding: connected Bytes Rx: 0 Tx: 0

```

show tcp-forwarding status (Listening Port and Address)

user@host> **show tcp-forwarding status listening-port 203.0.113.50 listening-address 8002**

```

Listening on: [default:]203.0.113.50:8002
  Status: listening
  Total Bytes Rx: 0 Tx: 0
Forwarding to: [default:]192.0.0.4:830
  Total Bytes Rx: 0 Tx: 0
Allowed Source Prefixes:
  198.51.100.3/32
Connections Max: 1 Active: 1
Source: 198.51.100.3:55002
  Listening: connected Bytes Rx: 0 Tx: 0
  Forwarding: connected Bytes Rx: 0 Tx: 0

```